arXiv:1402.1587v1 [cs.DM] 7 Feb 2014

Independent Set Reconfiguration in Cographs

Paul Bonsma*

February 10, 2014

Abstract

We study the following independent set reconfiguration problem, called *TAR-Reachability:* given two independent sets I and J of a graph G, both of size at least k, is it possible to transform I into J by adding and removing vertices one-by-one, while maintaining an independent set of size at least k throughout? This problem is known to be PSPACE-hard in general. For the case that G is a cograph (i.e. P_4-free graph) on n vertices, we show that it can be solved in time $O(n^2)$, and that the length of a shortest reconfiguration sequence from I to J is bounded by $4n - 2k$, if such a sequence exists.

More generally, we show that if $\mathcal{G}$ is a graph class for which (i) TAR-Reachability can be solved efficiently, (ii) maximum independent sets can be computed efficiently, and which satisfies a certain additional property, then the problem can be solved efficiently for any graph that can be obtained from a collection of graphs in $\mathcal{G}$ using disjoint union and complete join operations. Chordal graphs are given as an example of such a class $\mathcal{G}$.

1 Introduction

Reconfiguration problems have been studied often in recent years. These arise in settings where the goal is to transform feasible solutions to a problem in a step-by-step manner, while maintaining a feasible solution throughout. A *reconfiguration problem* is obtained by defining *feasible solutions* (or configurations) for *instances* of the problem, and a (symmetric) *adjacency relation* between solutions. This defines a *solution graph* for every instance, which is usually exponentially large in the input size. Usually, it is assumed that *adjacency* and *being a feasible solution* can be tested in polynomial time. Typical questions that are studied are deciding the existence of a path between two given solutions *(reachability)*, finding shortest paths between solutions, deciding whether the solution graph is connected or giving sufficient conditions for this, and giving bounds on its diameter. For example, the literature contains such results on the reconfiguration of vertex colorings [4, 6, 7, 8], boolean assignments that satisfy a given formula [15], independent sets [16, 19, 21, 22], matchings [19], shortest paths [2, 3, 20], subsets of a (multi-)set of integers [12, 18], etc. Techniques for many different reconfiguration problems are discussed in [19, 22]. See the recent survey by van den Heuvel [17] for an overview of and introduction to reconfiguration problems, and a discussion of their various applications.

One of the most well-studied problems of this kind is the reconfiguration of *independent sets*. For a graph G and integer k, the independent sets of size at least/exactly k of G form the feasible solutions. Independent sets are also called *token configurations*, where the

*University of Twente, Faculty of EEMCS, PO Box 217, 7500 AE Enschede, the Netherlands.

independent set vertices are viewed as *tokens*. Three types of adjacency relations have been studied in the literature: in the *token jumping (TJ)* model [19], a token can be moved from any vertex to any other vertex. In the *token sliding (TS)* model, tokens can be moved along edges of the graph [16]. In the *token addition and removel (TAR)* model [19], tokens can be removed and added in arbitrary order, though at least k tokens should remain at any time (k is the *token lower bound*). Of course, in all of these cases, an independent set should be maintained, so tokens can only be moved/added to vertices that are not dominated by the current token configuration.

The *reachability problem* has received the most attention in this context: given two independent sets I and J of a graph G, and possibly a token lower bound $k \leq \min\{|I|, |J|\}$, is there a path (or *reconfiguration sequence*) from I to J in the solution graph? We call this problem *TJ-Reachability, TS-Reachability* or *TAR-Reachability*, depending on the adjacency relation that is used. Kamiński et al [21] showed that the TAR-Reachability problem generalizes the TJ-Reachability problem (see Section 2 for details). For all three adjacency relations, this problem is PSPACE-hard, even in perfect graphs [21], and even in planar graphs of maximum degree 3 [16]. (The latter result is not explicitly stated in [16], but can easily be deduced from the given reduction. See [4] for more information.) See also [19] for an alternative, simple PSPACE-hardness proof. In addition, in [21], the problem of deciding whether there exists a path of length at most l between two solutions is shown to be strongly NP-hard, for all three adjacency models.

On the positive side, these problems can be solved in polynomial time for various restricted graph classes. The result on matching reconfiguration by Ito et al [19] implies that for line graphs, TJ-Reachability and TAR-Reachability can be solved efficiently. In [21], an efficient algorithm is given for TS-Reachability in cographs, and it is shown that for TJ-Reachability in even-hole-free graphs, a reconfiguration sequence exists between any pair of independent sets I and J, and that the shortest reconfiguration sequence always has length $|I \backslash J|$.

New results and techniques In this paper, we show that TAR-Reachability and TJ-Reachability can be solved in time $O(n^2)$ for cographs, where n is the number of vertices of the input graph. This answers an open question from [21]. In addition, we show that for cographs, components of the solution graph have diameter at most $4n - 2k$ and $2n - k$, under the TAR-model and TJ-model, respectively. Recall that a graph is a *cograph* iff it has no induced path on four vertices. Alternatively, cographs can be defined as graphs that can be obtained from a collection of trivial (one vertex) graphs by repeatedly applying *(disjoint) union* and *(complete) join* operations. The order of these operations can be described using a rooted *cotree*. This characterization allows efficient dynamic programming (DP) algorithms for various NP-hard problems. Our algorithm is also a DP algorithm over the cotree, albeit more complex than many known DP algorithms on cographs. For both solutions A and B, certain values are computed, using first a *bottom up* DP phase, and next a *top down* DP phase over the cotree. Using these values, we can conclude whether B is reachable from A. Because of this method, we in fact obtain a stronger result: TJ- and TAR-Reachability can be decided efficiently for any graph that can be obtained using join and union operations, when starting with a collection of base graphs from a graph class $\mathcal{G}$ that satisfies the following properties:

- For any graph in $\mathcal{G}$, the TAR-Reachability problem can be decided efficiently, and
- for any graph in $\mathcal{G}$ and independent set I, the size of a maximum independent set that is TAR-reachable from I can be computed efficiently, for all token lower bounds $k \leq |I|$.

In this paper, we show that an example of such a graph class is the class of *chordal graphs*. In another paper, we show that the class of *claw-free graphs* also satisfies these properties [5]. Combining these results yields quite a rich graph class for which this PSPACE-hard problem can be solved in polynomial time.

Another motivation for this research is that cographs form the base class for various graph width measures: cographs are exactly the graphs of cliquewidth at most two, and exactly the graphs of modular-width two [11]. The corresponding graph decompositions (k-expressions and modular decompositions) have been well-studied in algorithmic graph theory, because of the fact that many NP-hard problems can be solved efficiently on graphs where the width of these decompositions is low, using DP algorithms [10, 13]. Another similar, successful and widely used notion is that of a tree decomposition / the treewidth of a graph [1]. The success of such approaches for NP-complete problems and NP-optimization problems is unmistakable in the area of algorithmic graph theory. However, surprisingly, no nontrivial results of this kind are known for reconfiguration problems, to our knowledge. More precisely: we are not aware of any reconfiguration problems that are PSPACE-hard in general, but that can be solved efficiently on graphs of treewidth or cliquewidth at most k, for every constant k. On the other hand, none of the studied reconfiguration problems have been shown to be PSPACE-hard on graphs of bounded treewidth/cliquewidth. We expect that positive results of this kind are certainly possible, but have not yet been obtained due to the lack of DP techniques for reconfiguration problems. This paper gives a first example of how dynamic programming over graph decompositions can be used successfully for PSPACE-hard reconfiguration problems. This is a first step towards solving various reconfiguration problems for graphs of bounded (modular-, clique-, tree-) width; we expect that similar algorithmic techniques can be used and are necessary to show that indeed, various reconfiguration problems can be solved efficiently using DP over graph decompositions. We remark that a DP approach has also been used to show that the PSPACE-hard Shortest Path Reconfiguration problem can be solved in polynomial time on planar graphs [2], although a problem-specific layer decomposition of the graph was used.

Our DP algorithm for the TAR-Reachability problem is presented in Sections 4–6. First, in Section 3, an example is given, the proof of this statement is outlined, and a detailed overview of Sections 4–6 is given. In Section 7, examples of graph classes are given for which this algorithm works; in particular graphs obtained from chordal graphs using union and join operations (which includes cographs). The bound on the diameter of the solution graph is given in Section 8. We start in Section 2 with precise definitions, and end in Section 9 with a discussion.

2 Preliminaries

Token Addition and Removal By $\alpha(G)$ we denote the maximum size of an independent set in G. In this paper, we use the *token addition and removal (TAR)* model for independent set reconfiguration. For a graph G and integer k, the vertex set of the graph $\mathrm{TAR}_k(G)$ is the set of all independent sets of size at least k in G. Two distinct independent sets I and J are adjacent in $\mathrm{TAR}_k(G)$ if there exists a vertex $v \in V(G)$ such that $I \cup \{v\} = J$ or $I = J \cup \{v\}$. Vertices from independent sets will also be called *tokens*, and we will also say that J is obtained from I by *adding one token on* v resp. *removing one token from* v, or that J *is obtained from* I *using one TAR-step.*

For an integer k and two independent sets I and J of G with $|I| \geq k$ and $|J| \geq k$, we write $I \leftrightarrow_k^G J$ if $\mathrm{TAR}_k(G)$ contains a path from I to J. Observe that $I \leftrightarrow_0^G J$ always holds, and that the relation $\leftrightarrow_k^G$ is an equivalence relation, for all G and k. The superscript G is omitted if the graph in question is clear. If G and k are clear from the context, we will also simply say that J *is reachable from* I. A sequence $I_0, \ldots, I_k$ is called a *k-TAR-sequence* for G from I_0 to I_k if

- for every i, I_i is an independent set of G,
- for every i, $|I_i| \geq k$, and
- for every i, I_{i+1} can be obtained from I_i using at most one TAR-step.

Observe that $I \leftrightarrow_k^G J$ if and only if there exists a k-TAR-sequence in G from I to J. Note that we allow that $I_i = I_{i+1}$, in order to avoid discussing trivial cases in our proofs.

Our results also apply to the *token jumping (TJ)* model: for a graph G and integer k, the vertex set of the graph $\mathrm{TJ}_k(G)$ is the set of all independent sets of size *exactly* k in G. Two distinct independent sets I and J are adjacent in $\mathrm{TJ}_k(G)$ if there exist vertices $u \in I$ and $v \in J$ such that $I \backslash \{u\} = J \backslash \{v\}$. We say that J is obtained from I by *jumping a token from u to v*. Analogously to before, this defines TJ-sequences from I to J, and we write $I \leftrightarrow_{\mathrm{TJ}}^G J$ if a TJ-sequence from I to J exists. Kamiński et al showed that the TAR-model generalizes the TJ-model, in the following way:

Lemma 1 ([21]) *Let A and B be two independent sets of a graph G, with $|A| = |B| = \ell$. Then for any $k \in \mathbb{N}$, there exists an $(\ell - 1)$-TAR-sequence from A to B of length at most $2k$ if and only if there exists a TJ-sequence from A to B of length at most k.*

We remark that the TAR-model as defined in [21] is a little more restricted: for our algorithms, it is essential to consider the case where the token lower bound k is equal to the size of the initial independent sets A and B, whereas in [21], only the case where $k < \min\{|A|, |B|\}$ is considered.

Cographs and cotree decompositions For an illustration of the following definitions, see Figure 1. A *generalized cotree* is a binary tree T with root r, together with

- a partition of the nonleaf vertices into *union nodes* and *join nodes*, and
- a graph G_u for every leaf u of T, such that for any two leaves u and v, the graphs G_u and G_v are vertex and edge disjoint.

Vertices of T are called *nodes*. For every nonleaf node u, the two children are ordered; they are called the *left child* and *right child* of u. With every node $u \in V(T)$ we associate a graph G_u in the following way: for leaves u, G_u is as given. Otherwise, u has two child nodes; denote these by v and w. If u is a union node, then G_u is the *disjoint union* of G_v and G_w. If u is a join node, then G_u is obtained by taking the *complete join* of G_v and G_w. This operation is defined as follows: start with the disjoint union of G_v and G_w, and add edges yz for every combination of $y \in V(G_v)$ and $z \in V(G_w)$. For a node $u \in V(T)$, we denote $V_u = V(G_u)$. A generalized cotree T is called a *cotree* if for every leaf $v \in V(T)$, the graph G_v consists of a single vertex. Such a leaf is called a *trivial leaf*.

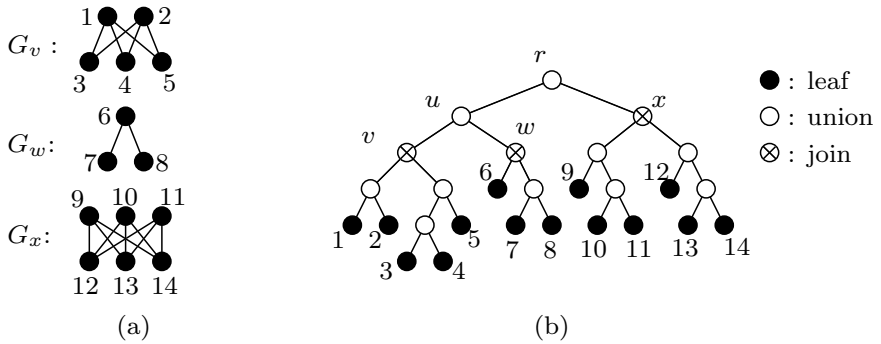

Figure 1: (a) A cograph G with components G_v, G_w and G_x, and (b) a cotree T of G. Leaves of T are labeled with the corresponding vertex number of G.

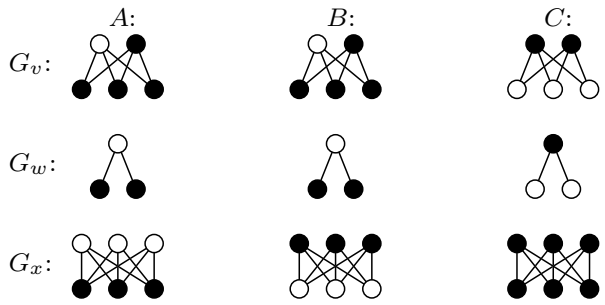

Figure 2: A cograph G, with independent sets A, B and C indicated by the white vertices. Any 5-TAR-sequence from A to B must visit C, use all vertices of G, and has length at least 24.

Let T be a (generalized) cotree, with root r. For a graph G, we say that T is a *(generalized) cotree for* G if $G_r = G$. A graph G is called a *cograph* if there exists a cotree for G. Let $\mathcal{G}$ be a graph class. We say that a generalized cotree T for a graph G is a *cotree decomposition of* G *into* $\mathcal{G}$*-graphs* if for every leaf $v \in V(T)$, the graph $G_v \in \mathcal{G}$. For instance, we will consider cotree decompositions into chordal graphs.

3 Example and Proof Outline

In this section, we will give an example, and use it to introduce the techniques and notions that will be used in the proofs. We will end with an outline of the algorithm, and overview of the paper.

Example In Figure 1, a cograph G together with a cotree T of G is shown. The root of T is r, and $V(G) = \{1, \ldots, 14\}$. The graph G has three components, which are G_v, G_w and G_x.

In Figure 2, three independent sets A, B and C are shown for the cograph G from Figure 1. In order to go from A to B in $\mathrm{TAR}_5(G)$, an independent set must be visited which has no tokens on the component G_x, and therefore at least five tokens on the other two components. The only such independent set of G is C. Using similar observations, it can be seen that there

the *shortest* 5-TAR-sequence from A (or B) to C is unique up to symmetries, and has length twelve (six additions and six deletions). Hence the shortest 5-TAR-sequence from A to B has length 24.

Proof Outline and Definitions For two independent sets A and B of a graph G, both with size at least k, we will characterize whether $A \leftrightarrow_k^G B$, using a (generalized) cotree for G. This requires the following notion.

Definition 2 *Let T be a generalized cotree for a graph G, I be an independent set of G, and $k \leq |I|$. For $v \in V(T)$, define $\lambda_k^I(v) = \min |J \cap V_v|$ over all independent sets J of G with $I \leftrightarrow_k^G J$.*

For instance, in the example from Figure 2, $\lambda_5^A(x) = 0 = \lambda_5^B(x)$, and this fact is essential for concluding that $A \leftrightarrow_5^G B$ in this case. In general, the following theorem characterizes whether B is reachable from A, using the values from Definition 2.

Theorem 3 *Let T be a generalized cotree for a graph G. Let A and B be two independent sets of G of size at least k. Then $A \leftrightarrow_k^G B$ if and only if*

1. *for all nodes $u \in V(T)$, $\lambda_k^A(u) = \lambda_k^B(u)$, and*
2. *for all leaves $u \in V(T)$, $(A \cap V_u) \leftrightarrow_\ell^{G_u} (B \cap V_u)$, where $\ell = \lambda_k^A(u)$.*

The forward direction of the statement is straightforward: if $A \leftrightarrow_k^G B$, then since $\leftrightarrow_k^G$ is an equivalence relation, any independent set J is reachable from A if and only if it is reachable from B. It follows that $\lambda_k^A(v) = \lambda_k^B(v)$ for all $v \in V(T)$. The second property follows by restricting all independent sets in a k-TAR-sequence from A to B to the subgraph G_v for any leaf $v \in V(T)$. By definition, these all have size at least $\ell = \lambda_k^A(v)$, so this yields an ℓ-TAR-sequence from $A \cap V_v$ to $B \cap V_v$ for G_v. For more details, see Section 6, where Theorem 3 is proved.

In order to efficiently decide whether $A \leftrightarrow_k^G B$, it remains to compute the values $\lambda_k^I(v)$ for all $v \in V(T)$ and $I = A, B$. How this can be done is shown in Section 5.4.

In the example from Figure 2, it holds that $\lambda_5^A(x) = 0$. This is because on the subgraph G_u, which is the disjoint union of components G_v and G_w, it is possible to reconfigure from the initial independent set A to an independent set with at least five tokens on G_u, while keeping at least two tokens on G_u throughout. This indicates that in order to compute the values $\lambda_k^I(v)$, the following values must be computed, for different values of $\ell \in \{0, \ldots, k\}$.

Definition 4 *Let T be a cotree for G, and I be an independent set of G. For $v \in V(T)$ and $\ell \in \{0, \ldots, |I \cap V_v|\}$, denote by $\mu_\ell^I(v)$ the maximum of $|J|$ over all independent sets J of G_v with $(I \cap V_v) \leftrightarrow_\ell^{G_v} J$.*

Note that the value $\mu_\ell^I(v)$ depends only on the situation in the subgraph G_v; not on the entire graph. This is in contrast to the values $\lambda_k^I(v)$. Observe also that $\mu_0^I(v) = \alpha(G_v)$ (regardless of the choice of I).

It is not obvious how to compute the values $\mu_\ell^I(u)$. For the example from Figure 2, concluding that $\mu_2^A(u) = 5$ requires studying the following 2-TAR-sequence for G_u. We start with one token on both G_v and G_w. One token can be added on G_v. This allows removing

the token from G_w, and subsequently moving to a better configuration, with two tokens on G_w. This in turn allows removing all tokens from G_v, and subsequently moving to a better configuration, with three tokens on G_v. A sequence of this type is called a *cascading sequence*. Informally, in such a sequence, we have a join node u with children v and w, and alternatingly move between on one hand a large independent set on v and a small independent set on w and on the other hand a large independent set on w and small independent set on v. The goal is to obtain ever larger independent sets until no more improvements can be made. In Section 5.2, we will show how to compute the values $\mu_\ell^I(v)$. This is done by characterizing the outcome of such cascading sequences, using *maximum ℓ-stable tuples*.

The values $\mu_\ell^I(u)$ for a node u with children v and w can be computed using only the values $\mu_{\ell'}^I(v)$ and $\mu_{\ell'}^I(w)$ for different choices of ℓ'. Hence these values can be computed using a *bottom up* dynamic programming algorithm, which starts at the leaves of the cotree. Next, the rules from Section 5.4 for computing the values $\lambda_k^I(u)$ can be used. As indicated by their definitions, computing these values requires considering the entire graph. Therefore this must be done using a *top down* dynamic programming algorithm, which starts at the root node of T. Together with Theorem 3, this yields our algorithm for deciding whether $A \leftrightarrow_k^G B$. Our main algorithmic result is summarized in the next theorem, which is proved in Section 6.

Theorem 5 *Let T be a generalized cotree for a graph G on n vertices, let $k \in \mathbb{N}$ and let A and B be independent sets of G. If for every nontrivial leaf $v \in V(T)$ and relevant integer ℓ,*

- *the values $\mu_\ell^A(v)$ and $\mu_\ell^B(v)$ are known, and*
- *it is known whether $(A \cap V_v) \leftrightarrow_\ell^{G_v} (B \cap V_v)$,*

then in time $O(n^2)$ it can be decided whether $A \leftrightarrow_k^G B$.

In particular, Theorem 5 implies that for any two independent sets A and B for a *cograph* G, it can be decided in time $O(n^2)$ whether $A \leftrightarrow_k^G B$.

In Section 8, we will give an upper bound for the length of a shortest k-TAR-sequence between two independent sets A and B. The above example shows that to go from A to B, it may be necessary to put tokens on vertices that are neither in A nor in B. Nevertheless, we can show that for a commonly reachable independent set C, there exists a k-TAR-sequence from A (resp. B) to C that for every vertex $v \in V(G)$, adds a token on v at most once. This shows that there exists a k-TAR-sequence from A to B of length at most $4n - |A| - |B|$.

For all of our proofs, an essential fact is that for every node u, the vertex set V_u is a module of G. We will first give lemmas related to independent set reconfiguration and modules in Section 4.

4 Module Lemmas

A *module* of a graph G is a set $M \subseteq V(G)$ such that for every $v \in V(G) \backslash M$, either $M \subseteq N(v)$ or $M \cap N(v) = \emptyset$. In other words: for every pair $u, v \in M$, $N(u) \backslash M = N(v) \backslash M$. Note that we will also consider $V(G)$ to be a (trivial) module of G. We will often use the following simple property of cographs.

Proposition 6 *Let T be a cotree of G. Then for any $v \in V(T)$, V_v is a module of G.*

Modules are very useful for independent set reconfiguration, since to some extent, we can reconfigure within the module and outside of the module independently. The following two lemmas make this more precise, and present two useful properties for the proofs below. These two lemma proofs also introduce proof techniques related to TAR-sequences that will be used often below. Later, we will however not apply them in the same level of detail again.

Lemma 7 *Let M be a module of a graph G, let k and y be integers, and let A be an independent set of G, with $|A \cap M| \geq \max\{1, y\}$ and $|A| \geq k$. Denote $H = G[M]$. If there exists an independent set B of G with $A \leftrightarrow_k^G B$ and $|B \cap M| \leq y$, and if there exists an independent set C of H with $(A \cap M) \leftrightarrow_y^H C$, then there exists an independent set D of G with $A \leftrightarrow_k^G D$ and $D \cap M = C$.*

Proof: Denote $A_M = A \cap M$, and $A_{\overline{M}} = A \backslash M$. First consider the case that $|A_M| = y$. Informally, we can then simply apply the same vertex additions and removals from the y-TAR-sequence from A_M to C to the entire independent set A, and this way maintain an independent set throughout.

Formally, let $I_0, \ldots, I_p$ be an y-TAR-sequence for H from A_M to C. Define $I'_i = I_i \cup A_{\overline{M}}$ for all i. Then $I'_0, \ldots, I'_p$ is the desired k-TAR-sequence from A to an independent set D of G with $D \cap M = C$. Indeed,

- for every i, both I_i and $A_{\overline{M}}$ are independent sets. Since M is a module and $|A_M| \geq 1$, $A_{\overline{M}}$ contains no vertices that are adjacent to any vertex in M, so I'_i is again an independent set of G.
- Since $|A_M| = y$, we have $|A_{\overline{M}}| \geq k - y$. By definition, for every i it holds that $|I_i| \geq y$, and thus $|I'_i| \geq y + k - y \geq k$.
- Clearly, every I'_{i+1} can be obtained from I'_i using at most one TAR-step.

In the remaining case, we may assume that $|A_M| \geq y + 1$. Consider a *shortest* k-TAR-seq $S = J_0, \ldots, J_q$ from A to *any* independent set B of G with $|B \cap M| \leq y$. So for every i with $i < q$, $|J_i \cap M| \geq y + 1$, and $B = J_q$ is obtained from J_{q-1} by removing a vertex from M. Since M is a module and J_{q-1} is an independent set, this implies that no vertex in $B \backslash M$ is adjacent to any vertex in M. Denote $B_{\overline{M}} = B \backslash M$.

Informally, we can now reverse the TAR-sequence S, but ignore every token addition or removal on $V(G) \backslash M$. This yields a k-TAR-sequence for G, from B to $A_M \cup B_{\overline{M}}$. Since $|B_{\overline{M}}| \geq k - y$, we can now apply the token additions and removals from the TAR-sequence from A_M to C to this independent set, similar to above, and obtain the desired independent set $D = C \cup B_{\overline{M}}$. Combining these three k-TAR-sequences shows that there is a k-TAR-sequence from A to D. We now define this more precisely, and verify that these are indeed TAR-sequences.

For every i, denote $J'_i = (J_i \cap M) \cup B_{\overline{M}}$. Consider the sequence $S' = J'_q, \ldots, J'_0$. The argue that this is a k-TAR-sequence from B to $A_M \cup B_{\overline{M}}$:

- As observed above, no vertex in $B_{\overline{M}}$ is adjacent to any vertex in M. Hence for every i, J'_i is an independent set.
- Recall that for every i, $|J_i \cap M| \geq y$, and $|B_{\overline{M}}| \geq k - y$, so $|J'_i| \geq k$.

- Clearly, consecutive sets in the sequence can be obtained from each other by at most one TAR-step.

Analog to the first part of the proof, one can show that there exists a k-TAR-sequence S'' for G from $A_M \cup B_{\overline{M}}$ to $C \cup B_{\overline{M}}$. Combining the sequences S, S' and S'' shows that $A \leftrightarrow_k^G (C \cup B_{\overline{M}})$, which proves the statement. □

Using similar techniques, we can prove the next lemma. This applies to the case where one module M can be partitioned into two sets M_1 and M_2, with no edges between them, which therefore are also modules.

Lemma 8 *Let M be a module of a graph G, such that M can be partitioned into two sets M_1 and M_2 with no edges between M_1 and M_2. Let A be an independent set of G, let B_1 be an independent set of G with $A \leftrightarrow_k^G B_1$, that maximizes $|B_1 \cap M_1|$ among all such sets, and let B_2 be an independent set of G with $A \leftrightarrow_k^G B_2$.*

Then there exists an independent set C of G with $A \leftrightarrow_k^G C$ and $C \cap M_i = B_i \cap M_i$, for $i \in \{1, 2\}$.

Proof: If $B_1 \cap M_1 = \emptyset$, then by definition of B_1, it also holds that $B_2 \cap M_1 = \emptyset$, and therefore chosing $C = B_2$ proves the statement. So now suppose that $B_1 \cap M_1 \neq \emptyset$.

Since the relation $\leftrightarrow_k^G$ is an equivalence relation, we conclude that there exists a k-TAR-sequence $S = I_0, \ldots, I_p$ from B_1 to B_2. We will use S to show that there exists an independent set C of G with $B_1 \leftrightarrow_k^G C$ and $C \cap M_i = B_i \cap M_i$, for $i \in \{1, 2\}$. Combining this with $A \leftrightarrow_k^G B_1$ shows that also $A \leftrightarrow_k^G C$, which proves the statement.

First suppose that S contains an independent set that contains no vertices of M. Then let I_i be the first such independent set, so $i \geq 1$ and I_i is obtained from I_{i-1} by removing a token from M. Since M is a module and I_{i-1} is an independent set, I_i therefore contains no vertex that is adjacent to any vertex in M. We can then simply add the vertices in $B_1 \cap M_1$ and $B_2 \cap M_2$ to I_i in any order, to obtain the desired independent set C (recall that there are no edges between M_1 and M_2). Combining the TAR-sequences from A to B_1, from B_1 to I_i, and from I_i to C shows that $A \leftrightarrow_k^G C$.

So now we may assume that every independent set I_i in the sequence S contains at least one vertex of M. Then we modify S as follows: we ignore all token additions and removals on M_1. We argue that this is a k-TAR-sequence from B_1 to $C = (B_1 \cap M_1) \cup (B_2 \backslash M_1)$. More precisely, for every i define $I'_i = (B_1 \cap M_1) \cup (I_i \backslash M_1)$, and consider the sequence $S' = I'_0, \ldots, I'_p$. We argue that S' is a k-TAR-sequence for G:

- Suppose to the contrary that there exists an i such that I'_i is not an independent set. Let i be the minimum index with this property. Then I_i is obtained from I_{i-1} by adding a vertex y that is adjacent to some vertex in B_1. Because vertices in M_2 are not adjacent to vertices in M_1, it follows that $y \notin M$. Since M is a module, y is adjacent to every vertex in M. But I_i contains at least one vertex of M, contradicting that it is an independent set. We conclude that every I'_i is an independent set.

- Let $y = |B_1 \cap M_1|$. Since every independent set I_i is also reachable from A, from the definition of B_1 it follows that $|I_i \cap M_1| \leq y$. Therefore, $|I_i \backslash M_1| \geq k - y$. It follows that for every i, $|I'_i| \geq y + k - y = k$.

- Clearly, consecutive sets in the sequence S' can be obtained from each other using at most one TAR-step.

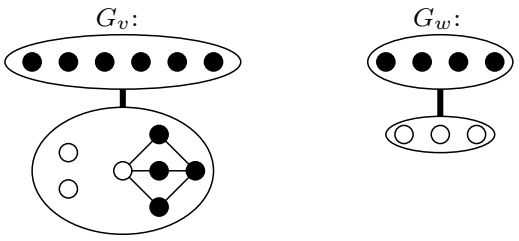

Figure 3: A cograph G_u with components G_v and G_w, and independent set I consisting of the white vertices.

We conclude that $B_1 \leftrightarrow_k^G C$. Combined with $A \leftrightarrow_k^G B_1$, it follows that $A \leftrightarrow_k^G C$. □

5 Dynamic Programming Rules

5.1 Cascading Sequences

In Section 5.2 below we will give dynamic programming rules for computing the values $\mu_\ell^I(u)$ for all nodes $u \in V(T)$. Recall that $\mu_\ell^I(u) = \max |J|$ where the maximum is taken over all independent sets J of G_u with $(I \cap V_u) \leftrightarrow_\ell^{G_u} J$. For trivial leaves and join nodes, the rules are straightforward. As discussed in Section 3, the computation for union nodes is more complicated, and requires studying the outcome of certain ℓ-TAR-sequences in G_u, which we will informally call *cascading sequences*.

We first introduce these informally, using the example shown in Figure 3. This figure depicts a cograph G_u which is obtained by taking the disjoint union of two cographs G_v and G_w. In this figure, a bold line between two encircled sets V_1 and V_2 of vertices means that edges xy are present between every $x \in V_1$ and $y \in V_2$. This corresponds to a complete join of $G[V_1]$ and $G[V_2]$. Let I be the independent set of G_u consisting of the white vertices. In Table 1, the values $\mu_\ell^I(v)$ and $\mu_\ell^I(w)$ are given for every $\ell \in \{0, \ldots, 3\}$. (These values can easily be verified. See also Figure 4 for examples of maximum independent sets that are reachable from $I \cap V_x$ in $\mathrm{TAR}_\ell(G_x)$ for various values of ℓ, and $x \in \{v, w\}$.)

Let the *type* of an independent set J of G_u be $(|J \cap V_v|, |J \cap V_w|)$. If it is required to keep at least $\ell = 6$ tokens on G_u throughout, then from the initial independent set I, which is of type $(3, 3)$, we can go to an independent set of type $(\mu_3^I(v), \mu_3^I(w)) = (4, 3)$. This holds by definition of $\mu_3^I(v)$ and $\mu_3^I(w)$, and because we can reconfigure in both components independently, as long as at least three tokens remain on both sides. From this, we could go to a independent set of type $(4, 2)$, but this does not enable further improvements. So we conclude that $\mu_6^I(u) = 4 + 3 = 7$. If $\ell = 5$, then observe that we can go from the initial independent set of type $(3, 3)$ to one of type $(2, 3)$, and subsequently to one of type $(\mu_2^I(v), 3) = (5, 3)$. Next, we can visit independent sets of types $(5, 0)$ and $(5, \mu_0^I(w)) = (5, 4)$. We could then go to one of type $(1, 4)$, but since $\mu_1^I(v) = 5$, this yields no improvement. We conclude that $\mu_5^I(u) = 5 + 4 = 9$. Finally, if $\ell \leq 4$, then we can visit independent sets of types $(3, 3)$, $(4, 3)$, $(4, 0)$, $(4, 4)$, $(0, 4)$, $(6, 4)$, in this order, using similar arguments. Since $6 + 4 = 10 = \alpha(G_u)$, no further improvements are possible, so $\mu_\ell^I(u) = 10$ for all $\ell \leq 4$. This yields the values $\mu_\ell^I(u)$ shown in Table 1. Note that we can deduce these values using only the previous two columns of the table; without considering other details about the graph.

ℓ :	$\mu_\ell^I(v)$:	$\mu_\ell^I(w)$:	$\mu_\ell^I(u)$:
0	6	4	10
1	5	3	10
2	5	3	10
3	4	3	10
4			10
5			9
6			7

Table 1: The values $\mu_\ell^I(x)$ for $x \in \{u, v, w\}$ and $\ell \in \{0, \ldots, |I \cap V_x|\}$.

Below we will prove that the values computed this way are correct. However, we will not formalize cascading sequences, but instead characterize their outcome. We will define *maximum ℓ-stable tuples* (x, y) for each $\ell \in \{0, \ldots, |I \cap V_u|\}$, and show that $x = \min |J \cap V_v|$ and $y = \min |J \cap V_w|$, where in both cases the minimum is taken over all independent sets J of G_u with $(I \cap V_u) \leftrightarrow_\ell^{G_u} J$. So for the example above and $\ell = 6, 5, 4$, these tuples can be verified to be $(3, 2)$, $(1, 0)$ and $(0, 0)$ respectively. (As indicated by the above cascading sequences.) Next, we will show that $\mu_\ell^I(u) = \mu_x^I(v) + \mu_y^I(w)$, where (x, y) is the maximum ℓ-stable tuple. The maximum ℓ-stable tuple can easily be computed from its definition, given below.

5.2 Bottom Up Dynamic Programming Rules

Throughout this section, T denotes a generalized cotree of G and I denotes an independent set of G. The following property follows easily from the definition of $\mu_\ell^I(u)$, and will often be used in this section.

Proposition 9 *Let $u \in V(T)$. For any two integers x, y with $0 \leq x \leq y \leq |I \cap V_u|$, it holds that $\mu_x^I(u) \geq \mu_y^I(u)$.*

For trivial leaf nodes, the computation of these values is easy:

Proposition 10 *Let $u \in V(T)$ be a trivial leaf node. Then $\mu_\ell^I(u) = 1$ for all ℓ.*

For join nodes, the computation of $\mu_\ell^I(u)$ is still relatively straightforward. Note that for any join node u and independent set I, u has a child w with $V_w \cap I = \emptyset$.

Proposition 11 *Let $u \in V(T)$ be a join node. Let w be a child of u with $I \cap V(G_v) = \emptyset$, and let v be the other child of u.*

- *$\mu_\ell^I(u) = \mu_\ell^I(v)$ for all $\ell \geq 1$, and*
- *$\mu_0^I(u) = \max\{\mu_0^I(v), \mu_0^I(w)\}$.*

Proof: Because all edges are present between G_v and G_w, a maximum independent set of G_u is either a maximum independent set of G_v or of G_w, so $\mu_0^I(u) = \alpha(G_u) = \max\{\alpha(G_v), \alpha(G_w)\} = \max\{\mu_0^I(v), \mu_0^I(w)\}$. Now consider the case $\ell \geq 1$, and thus $|I \cap V_u| \geq 1$. Then initially all

tokens of I are on the child G_v. As long as there is at least one token on G_v, no tokens can be added to G_w. So essentially, G_w can be ignored, and thus $\mu_\ell^I(u) = \mu_\ell^I(v)$. □

For union nodes u, we will show that the value of $\mu_\ell^I(u)$ can be characterized using *maximum stable tuples*, which are defined as follows.

Definition 12 *For a union node u with left child v and right child w, independent set $I \subseteq V(G)$ and integer $\ell \leq |I \cap V_u|$, call a tuple (x, y) of integers with $x \leq |I \cap V_v|$ and $y \leq |I \cap V_w|$* ℓ-stable *if*

$$x = \max\{0, \ell - \mu_y^I(w)\} \text{ and } y = \max\{0, \ell - \mu_x^I(v)\}.$$

Call an ℓ-stable tuple (x, y) maximum *if there is no ℓ-stable tuple (x', y') with $x' \geq x$, $y' \geq y$ and $(x, y) \neq (x', y')$.*

In the remainder of this section, we will first prove that for every ℓ, there exists a unique maximum ℓ-stable tuple (x, y), and characterize this tuple (Lemma 13 below). Using this characterization, we can show that for a union node u with children v and w and any ℓ, $\mu_\ell^I(u) = \mu_x^I(v) + \mu_y^I(w)$, where (x, y) is the unique maximum ℓ-stable tuple (Lemma 19 below).

Lemma 13 *Let $u \in V(T)$ be a union node, with left child v and right child w. For $\ell \in \{0, \ldots, |I \cap V_u|\}$, let $x = \min |J \cap V_v|$ and $y = \min |J \cap V_w|$, where in both cases the minimum is taken over all independent sets J of G_u with $(I \cap V_u) \leftrightarrow_\ell^{G_u} J$. Then (x, y) is the unique maximum ℓ-stable tuple for I and u.*

Before we can prove Lemma 13, we first need to prove a number of other statements. These statements will refer to notations I, u, v, w, x, y, ℓ as defined in Lemma 13. In addition, we will denote $I_u = I \cap V_u$.

Proposition 14 *Consider a TAR-sequence $S = J_0, \ldots, J_p$ in $TAR_\ell(G_u)$ with $J_0 = I_u$. Let $x^* = \min_i |J_i \cap V_v|$ and $y^* = \min_i |J_i \cap V_w|$. Then $\mu_{x^*}^I(v) \geq |J_p \cap V_v|$ and $\mu_{y^*}^I(w) \geq |J_p \cap V_w|$.*

Proof: Consider the sequence $S' = J'_0, \ldots, J'_p$ with $J'_i = J_i \cap V_w$ for all i. This is a y^*-TAR-sequence for G_w that ends with J'_p. So by definition, $\mu_{y^*}^I(w) \geq |J'_p|$. The proof of the other statement is analog. □

Since x and y (as defined in Lemma 13) provide lower bounds for x^* and y^* respectively (as defined in Proposition 14), we conclude:

Corollary 15 *For every J with $I_u \leftrightarrow_\ell^{G_u} J$, it holds that $\mu_x^I(v) \geq |J \cap V_v|$ and $\mu_y^I(w) \geq |J \cap V_w|$.*

Using Proposition 14, we can draw the following two conclusions.

Proposition 16 *$x \geq \ell - \mu_y^I(w)$ and $y \geq \ell - \mu_x^I(v)$.*

Proof: Consider an ℓ-TAR-sequence $J_0, \ldots, J_p$ for G_u from I_u to an independent set J_p with $|J_p \cap V_v| = x$. Let $y^* = \min_i |J_i \cap V_w|$, so $y^* \geq y$. By Proposition 14, $\mu_{y^*}^I(w) \geq |J_p \cap V_w| \geq \ell - |J_p \cap V_v| = \ell - x$. Using Proposition 9 and $y^* \geq y$, it follows that $\mu_y^I(w) \geq \ell - x$ holds as well. The other inequality is proved analogously. □

Lemma 17 *For any ℓ-stable tuple (x', y'), it holds that $x \geq x'$ and $y \geq y'$.*

Proof: Suppose to the contrary that there exists an ℓ-TAR-sequence for G_u from I_u to some independent set J with $|J \cap V_v| < x'$ or $|J \cap V_w| < y'$. Consider a *shortest* ℓ-TAR-sequence $S = J_0, \ldots, J_p$ of this kind, and assume w.l.o.g. this ends with J_p with $|J_p \cap V_v| = x' - 1$. This implies that $x' \geq 1$, and therefore $x' = \ell - \mu^I_{y'}(w)$ (since (x', y') is stable). It follows that $|J_p \cap V_w| \geq \ell - |J_p \cap V_v| = \ell - x' + 1 = \mu^I_{y'}(w) + 1$, so

$$|J_p \cap V_w| \geq \mu^I_{y'}(w) + 1. \tag{1}$$

Combining this with the trivial lower bounds $\mu^I_{y'}(w) \geq |I \cap V_w| \geq y'$ we obtain

$$|J_p \cap V_w| \geq y' + 1.$$

Let $y^* = \min_i |J_i \cap V_w|$, and choose i accordingly such that $|J_i \cap V_w| = y^*$. Combining Proposition 14 with (1) yields

$$\mu^I_{y^*}(w) \geq |J_p \cap V_w| \geq \mu^I_{y'}(w) + 1.$$

It follows that $y^* < y'$ (Proposition 9). So $|J_p \cap V_w| \geq y' + 1 > y^* + 1 = |J_i \cap V_w| + 1$, and thus $i < p$. But then the subsequence of S that ends with J_i satisfies $|J_i \cap V_w| = y^* < y'$, and this is a strictly shorter sequence than S, a contradiction with the choice of S. □

Proposition 18 *There exists an independent set J_1 of G_u with $|J_1 \cap V_v| = \mu^I_x(v)$ and $I_u \leftrightarrow^{G_u}_\ell J_1$, and there exists an independent set J_2 of G_u with $|J_2 \cap V_w| = \mu^I_y(w)$ and $I_u \leftrightarrow^{G_u}_\ell J_2$.*

Proof: We prove the second statement. The proof of the first statement is analog. If $I_u \cap V_w = \emptyset$, then we can simply add vertices from a maximum independent set C of G_w to I_u, one by one. Recall that $|C| = \alpha(G_w) = \mu^I_0(w)$. Since G_u is the disjoint union of G_v and G_w, this yields a TAR-sequence in G_u, from I_u to a independent set J_2 with $|J_2 \cap V_w| = \mu^I_0(w) = \mu^I_y(w)$.

So we may now assume that $|I_u \cap V_w| \geq 1$, and we can apply (module) Lemma 7, with I_u in the role of A, V_w in the role of the module M and G_u in the role of the entire graph G. By definition of y, there exists an independent set B of G_u with $I_u \leftrightarrow^{G_u}_\ell B$ and $|B \cap V_w| = y$. By definition of $\mu^I_y(w)$, there exists an independent set C of G_w with $(I_u \cap V_w) \leftrightarrow^{G_w}_y C$ and $|C| = \mu^I_y(w)$. Now Lemma 7 shows that there exists an independent set J_2 of G_u with $A \leftrightarrow^{G_u}_\ell J_2$ and $|J_2 \cap V_w| = |C| = \mu^I_y(w)$. □

Now we are ready to prove Lemma 13.

Proof of Lemma 13: Consider J_2 as in Proposition 18. We can remove all but $\max\{0, \ell - \mu^I_y(w)\}$ tokens from G_v, and still have at least ℓ tokens in total on G_u. This shows that $x \leq \max\{0, \ell - \mu^I_y(w)\}$. Analogously, $y \leq \max\{0, \ell - \mu^I_x(v)\}$ follows. Combining these inequalities with Proposition 16 and the obvious inequalities $x \geq 0$, $y \geq 0$ shows that $x = \max\{0, \ell - \mu^I_y(w)\}$ and $y = \max\{0, \ell - \mu^I_x(v)\}$, hence the tuple (x, y) is ℓ-stable. Furthermore, Lemma 17 shows that (x, y) is a *maximum* ℓ-stable tuple, and in fact the *only* maximum ℓ-stable tuple. □

Lemma 13 implies in particular that there exists a unique maximum ℓ-stable tuple for any choice of ℓ. From now on we will now use this fact implicitly, for instance in the following lemma statement. Now we are ready to state and prove Lemma 19, which shows how the values $\mu^I_\ell(u)$ can be computed for a join node u.

Lemma 19 *Let $u \in V(T)$ be a union node, with left child v and right child w. For $\ell \in \{0, \ldots, |I \cap V_u|\}$, let (x, y) be the unique maximum ℓ-stable tuple for I and u. Then $\mu^I_\ell(u) = \mu^I_x(v) + \mu^I_y(w)$.*

Proof: Lemma 13 shows that $x = \min |J \cap V_v|$ and $y = \min |J \cap V_w|$, where in both cases the minimum is taken over all independent sets J of G_u with $(I \cap V_u) \leftrightarrow^{G_u}_\ell J$, so we may apply the above statements that were proved for this choice of x and y.

For any independent set J^* of G_u with $I_u \leftrightarrow^{G_u}_\ell J^*$, Corollary 15 shows that $|J^*| = |J^* \cap V_v| + |J^* \cap V_w| \leq \mu^I_x(v) + \mu^I_y(w)$, so

$$\mu^I_\ell(u) \leq \mu^I_x(v) + \mu^I_y(w). \tag{2}$$

Now it suffices to prove that

$$\mu^I_\ell(u) \geq \mu^I_x(v) + \mu^I_y(w). \tag{3}$$

To this end, we will show that (module) Lemma 8 can be applied, with G_u in the role of the entire graph G, V_u in the role of the module M, and V_v and V_w in the roles of the modules M_1 and M_2, respectively (recall that $\{M_1, M_2\}$ should be a partition of M with no edges between M_1 and M_2). We choose I_u in the role of A. By Proposition 18, there exists an independent set J_1 of G_u with $|J_1 \cap V_v| = \mu^I_x(v)$ and $I_u \leftrightarrow^{G_u}_\ell J_1$. Corollary 15 shows that J_1 maximizes the number of vertices on V_v among all reachable sets. Analogously, these two propositions show that there exists an independent set J_2 of G_u with $|J_2 \cap V_v| = \mu^I_y(w)$ and $I_u \leftrightarrow^{G_u}_\ell J_2$, that maximizes the number of vertices on V_w among all reachable independent sets. When using J_1 and J_2 in the roles of B_1 and B_2, Lemma 8 shows that there exists an independent set C of G with $I_u \leftrightarrow^{G_u}_\ell C$, $C \cap V_v = J_1 \cap V_v$ and $C \cap V_w = J_2 \cap V_w$. Inequality (3) follows since $|J_1 \cap V_v| = \mu^I_x(v)$ and $|J_2 \cap V_w| = \mu^I_y(w)$. □

5.3 Computing the Values Efficiently

Let u be a union node with children v and w such that for every relevant integer ℓ, the values $\mu^I_\ell(v)$ and $\mu^I_\ell(w)$ are known. Then Lemma 19 shows that for every relevant value ℓ, the value $\mu^I_\ell(u)$ can be computed in polynomial time: try all relevant combinations (x, y), verify whether they are stable, and subsequently identify the unique maximum stable tuple. However, this is not very efficient. In this section we present a more efficient method for computing the values $\mu^I_\ell(u)$ for union nodes u. The method is shown in Algorithm 1.

To prove that Algorithm 1 is correct, we need the following invariant.

Proposition 20 *At any time during the computation of Algorithm 1, for the variables a, b and ℓ the following property holds: For any $\ell' \in \{0, \ldots, \ell\}$ and ℓ'-stable tuple (x, y): $a \geq x$ and $b \geq y$.*

Proof: Consider the initial choices $a = |I \cap V_v|$, $b = |I \cap V_w|$ and $\ell = a + b$, and any ℓ'-stable tuple (x, y) for $\ell' \leq \ell$. If $x = 0$, then obviously $a \geq x$. Otherwise, $x = \ell' - \mu^I_y(w) \leq \ell - |I \cap V_w| = a$. Analogously, $b \geq y$ follows.

Now suppose that the claim holds for (a', b'), and that (a, b) is obtained from this tuple as shown in the algorithm. More precisely, (x, y) is an ℓ'-stable tuple for $\ell' \leq \ell$, and we have $a' \geq x$, $b' \geq y$, $a = \max\{0, \ell - \mu^I_{b'}(w)\}$ and $b = \max\{0, \ell - \mu^I_{a'}(v)\}$. We prove that $b \geq y$. The case $y = 0$ is trivial, so now assume $y \geq 1$, and therefore by ℓ'-stability of

Algorithm 1 Efficiently computing all values $\mu_\ell^I(u)$ for a union node u.

INPUT: For a union node u with children v and w: The values $|I \cap V_z|$ and $\mu_\ell^I(z)$ for $z \in \{v, w\}$ and $\ell \in \{0, \ldots, |I \cap V_z|\}$.
OUTPUT: Values $\mu_\ell^I(u)$ for all $\ell \in \{0, \ldots, |I \cap V_u|\}$.

```
 1: a := |I ∩ V_v|
 2: b := |I ∩ V_w|
 3: ℓ := a + b
 4: while ℓ ≥ 0 do
 5:     repeat
 6:         a' := a
 7:         b' := b
 8:         a := max{0, ℓ − μ^I_{b'}(w)}
 9:         b := max{0, ℓ − μ^I_{a'}(v)}
10:     until a = a' and b = b'
11:     μ^I_ℓ(u) := μ^I_a(v) + μ^I_b(w)
12:     ℓ := ℓ − 1
13: endwhile
```

(x, y), $y = \ell' - \mu_x^I(v)$. Since $a' \geq x$, Proposition 9 yields $\mu_{a'}^I(v) \leq \mu_x^I(v)$. We conclude that $b \geq \ell - \mu_{a'}^I(v) \geq \ell' - \mu_x^I(v) = y$. The inequality $a \geq x$ follows analogously. It follows that the assignments in Lines 8 and 9 maintain the invariant. Finally, decreasing ℓ by one (Line 12) also obviously maintains the invariant. □

Lemma 21 *Algorithm 1 correctly computes the values $\mu_\ell^I(u)$ for all $\ell \in \{0, \ldots, |I \cap V_u|\}$.*

Proof: The repeat-until loop terminates when $a = \max\{0, \ell - \mu_b^I(w)\}$ and $b = \max\{0, \ell - \mu_a^I(v)\}$, so when (a, b) is ℓ-stable for I and u. By Proposition 20, for any ℓ-stable tuple (x, y) it holds that $a \geq x$ and $b \geq y$, so (a, b) is a maximum ℓ-stable tuple. Then Lemma 19 shows that the assignment $\mu_\ell^I(u) = \mu_a^I(v) + \mu_b^I(w)$ is correct. □

To bound the complexity of Algorithm 1, we need the following invariant.

Proposition 22 *At any time during the computation of Algorithm 1, for the variables a, b and ℓ the following property holds: $a \geq \ell - \mu_b^I(w)$ and $b \geq \ell - \mu_a^I(v)$.*

Proof: For the initial choices of a, b and ℓ, the claim holds, since $\ell = a+b$, $\mu_b^I(w) \geq |I \cap V_w| = b$ and $\mu_a^I(v) \geq |I \cap V_v| = a$. Now suppose that the claim holds for (a', b'), and that (a, b) is obtained from this tuple as shown in the algorithm. More precisely, we have $a' \geq \ell - \mu_{b'}^I(w)$, $b' \geq \ell - \mu_{a'}^I(v)$, $a = \max\{0, \ell - \mu_{b'}^I(w)\}$ and $b = \max\{0, \ell - \mu_{a'}^I(v)\}$. We first argue that $a \leq a'$. If $a = 0$, the statement is clear (using the obvious invariant that these values remain nonnegative). Otherwise, we can write $a = \ell - \mu_{b'}^I(w) \leq a'$. By Proposition 9, it follows that $\mu_a^I(v) \geq \mu_{a'}^I(v)$, and therefore $b \geq \ell - \mu_{a'}^I(v) \geq \ell - \mu_a^I(v)$. Analogously, $a \geq \ell - \mu_b^I(w)$ follows. This shows that the assignments in Lines 8 and 9 maintain the invariant. Clearly, decreasing ℓ by one (Line 12) maintains the invariant as well. □

Lemma 23 *Algorithm 1 terminates in time $O(|I \cap V_u|)$.*

Proof: Denote $n = |I \cap V_u|$. All lines take constant time (we may assume that the quantities $|I \cap V_v|$ and $|I \cap V_w|$ are known). Therefore it suffices to show that in total, the variables a and b are reassigned at most $2n$ times (in Lines 8 and 9). Proposition 22 shows that whenever a new tuple (a, b) is obtained from a previous tuple (a', b'), that $a' \geq a$ and $b' \geq b$. (If $a = 0$, the statement is clear, and otherwise $a = \ell - \mu^I_{b'}(w) \leq a'$ holds. $b \leq b'$ follows similarly.) If $a' = a$ and $b' = b$, then ℓ is subsequently decreased (Line 12), which is done n times in total. Otherwise, $a' + b' > a + b$, and since both values remain nonnegative throughout, this can occur at most n times as well. Hence no line of the algorithm is visited more than $2n$ times. □

Theorem 24 *Let T be a generalized cotree of a graph G on n vertices, and let I be an independent set of G. If the values $\mu^I_\ell(v)$ are known for all nontrivial leaves $v \in V(T)$ and all relevant integers ℓ, then there is an algorithm that computes*

- *the values $\mu^I_\ell(u)$ for all $u \in V(T)$ and $\ell \in \{0, \ldots, |I \cap V_u|\}$, and*
- *the maximum ℓ-stable tuples for all union nodes $u \in V(T)$ and $\ell \in \{0, \ldots, |I \cap V_u|\}$,*

with time complexity $O(M) \subseteq O(n^2)$, where $M = \sum_{u \in V(T)} |I \cap V_u|$.

Proof: The Lemmas 21 and 23 show that for a union node u, Algorithm 1 computes the values $\mu^I_\ell(u)$ for all $\ell \in \{0, \ldots, |I \cap V_u|\}$ in time $O(|I \cap V_u|)$, given that the corresponding values for all child nodes are known. For the case that u is a trivial leaf or join node, the same claim follows easily from Propositions 10 and 11. So, using a straightforward bottom up computation, all values $\mu^I_\ell(u)$ can be computed correctly in time $O(M)$. That is, in constant time on average per entry. It remains to bound M in terms of n.

For a node $u \in V(T)$, define $f(u) = \sum_v |V_v|$, where the sum is over all descendants v of u in T, including u itself. By induction over T, we show that $f(u) \leq |V_u|^2$. The induction base is trivial. For the induction step, consider a node u with children v and w, and write $a = |V_v|$ and $b = |V_w|$. Then using the induction hypothesis, we can write

$$f(u) = |V_u| + f(v) + f(w) \leq (a + b) + a^2 + b^2 \leq 2ab + a^2 + b^2 = (a + b)^2 = |V_u|^2.$$

(We used $a \geq 1$ and $b \geq 1$.) Let r be the root of T. Using

$$M = \sum_{u \in V(T)} |I \cap V_u| \leq \sum_{u \in V(T)} |V_u| = f(r) \leq n^2,$$

the statement follows. □

5.4 Top Down Dynamic Programming Rules

Throughout this section, T denotes again a generalized cotree of G and I denotes an independent set of G. In this section, we will show how the values $\lambda^I_k(v)$ can be computed for all nodes $v \in V(T)$. For the case that v is a union node, this requires knowledge of a maximum ℓ-stable tuple (characterized in Lemma 13). For the root node of T, the value is trivial.

Proposition 25 *Let r be the root node of the cotree T. Then $\lambda_k^I(r) = k$.*

Proposition 26 *Let $u \in V(T)$ be a join node, with children v and w such that $I \cap V_w = \emptyset$. Then $\lambda_k^I(v) = \lambda_k^I(u)$ and $\lambda_k^I(w) = 0$.*

Proof: Considering I itself, $\lambda_k^I(w) = 0$ follows immediately. The inequality $\lambda_k^I(v) \leq \lambda_k^I(u)$ follows since $V_v \subseteq V_u$. It remains to prove that $\lambda_k^I(v) \geq \lambda_k^I(u)$.

Consider a *shortest* k-TAR-sequence $I_0, \ldots, I_p$ in G from I to any independent set I_p with $|I_p \cap V_v| = \lambda_k^I(v)$. If $p = 0$, then $\lambda_k^I(u) \leq |I \cap V_u| = |I \cap V_v| = \lambda_k^I(v)$, so now assume $p \geq 1$. Then I_p is obtained from I_{p-1} by removing a vertex from V_v. Since G_u is the complete join of G_v and G_w and I_{p-1} is an independent set, $I_{p-1} \cap V_w = \emptyset$. So $\lambda_k^I(u) \leq |I_p \cap V_u| = |I_p \cap V_v| = \lambda_k^I(v)$. □

Lemma 27 *Let $u \in V(T)$ be a union node, with left child v and right child w. Let $\ell = \lambda_k^I(u)$, and let (x, y) be the maximum ℓ-stable tuple for I and u. Then $\lambda_k^I(v) = x$ and $\lambda_k^I(w) = y$.*

Proof: Denote again $I_u = I \cap V(G_u)$. By Lemma 13, for the maximum ℓ-stable tuple (x, y) for I and u it holds that

$$x = \min |J \cap V_v| \text{ and } y = \min |J \cap V_w|, \tag{4}$$

where in both cases the minimum is taken over all independent sets J of G_u with $I_u \leftrightarrow_\ell^{G_u} J$.

We first use this to show that $\lambda_k^I(v) \geq x$ and $\lambda_k^I(w) \geq y$. Consider a k-TAR-sequence $I_0, \ldots, I_p$ for G with $I_0 = I$ and $|I_p \cap V_v| = \lambda_k^I(v)$. For every i, denote $I_i' = I_i \cap V_u$, and consider the sequence $I_0', \ldots, I_p'$. By definition of $\ell = \lambda_k^I(u)$, for every i it holds that $|I_i'| \geq \ell$, so this is an ℓ-TAR-sequence for G_u, and thus $I_u \leftrightarrow_\ell^{G_u} I_p'$. Using (4) it then follows that $\lambda_k^I(v) = |I_p \cap V_v| \geq x$. Analogously, $\lambda_k^I(w) \geq y$ follows.

We will now prove that $\lambda_k^I(v) \leq x$ and $\lambda_k^I(w) \leq y$. By (4), there exist independent sets J_1 and J_2 of G_u with $I_u \leftrightarrow_\ell^{G_u} J_1$, $I_u \leftrightarrow_\ell^{G_u} J_2$, $|J_1 \cap V_v| = x$ and $|J_2 \cap V_w| = y$. By the definition of $\ell = \lambda_k^I(u)$, there exists an independent set B of G with $I \leftrightarrow_k^G B$ and $|B \cap V_u| = \ell$. We can now apply (module) Lemma 7 twice, with V_u in the role of module M, I_u in the role of A, and J_1 or J_2 respectively in the role of C to conclude that there exist independent sets D_1 and D_2 of G with $I \leftrightarrow_k^G D_1$, $I \leftrightarrow_k^G D_2$, $|D_1 \cap V_v| = x$ and $|D_2 \cap V_w| = y$. Thus $\lambda_k^I(v) \leq x$ and $\lambda_k^I(w) \leq y$. □

6 Algorithm Summary and Main Theorems

In this section, we prove the two main theorems, and summarize how the previous facts and dynamic programming rules can be used to decide efficiently whether $A \leftrightarrow_k^G B$ for any two given independent sets A and B of a G, for any graph that satisfies the properties stated in Theorem 5. First, we prove the theorem that characterizes whether $A \leftrightarrow_k^G B$, using the previously defined values.

Theorem 3 *Let T be a generalized cotree for a graph G. Let A and B be two independent sets of G of size at least k. Then $A \leftrightarrow_k^G B$ if and only if*

1. *for all nodes* $u \in V(T)$, $\lambda_k^A(u) = \lambda_k^B(u)$, *and*

2. *for all leaves* $u \in V(T)$, $(A \cap V_u) \leftrightarrow_\ell^{G_u} (B \cap V_u)$, *where* $\ell = \lambda_k^A(u)$.

Proof: We first prove the forward direction. Suppose that $A \leftrightarrow_k^G B$. Then clearly, for any independent set J of G, $A \leftrightarrow_k^G J$ holds if and only if $B \leftrightarrow_k^G J$. So $\lambda_k^A(u) = \lambda_k^B(u)$ holds for every node $u \in V(T)$ (Definition 2), which proves the first property.

For any $u \in V(T)$, we may now denote $\lambda_k(u) = \lambda_k^A(u) = \lambda_k^B(u)$. Consider a k-TAR-sequence $I_0, \ldots, I_p$ for G from A to B. For any node $u \in V(T)$ and any $i \in \{0, \ldots, p\}$, it holds that $|I_i \cap V_u| \geq \lambda_k(u)$ (Definition 2). So $I'_0, \ldots, I'_p$ with $I'_i = I_i \cap V_u$ for all i is a $\lambda_k(u)$-TAR-sequence for G_u. This shows that $(A \cap V_u) \leftrightarrow_{\lambda_k(u)}^{G_u} (B \cap V_u)$, and thus proves the second property.

Now we prove the other direction. Assume that the two properties hold. So we may denote $\lambda_k(u) = \lambda_k^A(u) = \lambda_k^B(u)$ for all nodes u. We prove the following claim by induction over T:

Claim A: For all nodes $u \in V(T)$: $(A \cap V_u) \leftrightarrow_{\lambda_k(u)}^{G_u} (B \cap V_u)$.

Induction base: For leaf nodes $u \in V(T)$, the statement follows immediately from the second property.

Induction step: First consider a join node $u \in V(T)$ with children v and w. Suppose that $\lambda_k(v) \geq 1$. This implies $A \cap V_v \neq \emptyset$ and $B \cap V_v \neq \emptyset$. Therefore, since u is a join node, $A \cap V_u = A \cap V_v$ and $B \cap V_u = B \cap V_v$. In addition, $\lambda_k(u) = \lambda_k(v)$ (Proposition 26). From these facts, and the induction assumption $(A \cap V_v) \leftrightarrow_{\lambda_k(v)}^{G_v} (B \cap V_v)$, we conclude that $(A \cap V_u) \leftrightarrow_{\lambda_k(u)}^{G_u} (B \cap V_u)$. The case $\lambda_k(w) \geq 1$ is analog. Now suppose that $\lambda_k(v) = \lambda_k(w) = 0$. Then $\lambda_k(u) = 0$ (Proposition 26). The desired claim follows since $(A \cap V_u) \leftrightarrow_0^{G_u} (B \cap V_u)$ trivially holds.

Next, consider the case that $u \in V(T)$ is a union node with left child v and right child w. Denote $\ell = \lambda_k(u)$, $x = \lambda_k(v)$ and $y = \lambda_k(w)$. By Lemma 27, (x, y) is the maximum ℓ-stable tuple for u, for both A and B. We define C_v to be an independent set of G_v with $(A \cap V_v) \leftrightarrow_x^{G_v} C_v$, with maximum size among all such sets, and define C_w to be an independent set of G_w with $(A \cap V_w) \leftrightarrow_y^{G_w} C_w$, with maximum size among all such sets. By induction, $(A \cap V_v) \leftrightarrow_x^{G_v} (B \cap V_v)$, so it also holds that $(B \cap V_v) \leftrightarrow_x^{G_v} C_v$, and that C_v has maximum size among all such reachable sets. Analogously, $(B \cap V_w) \leftrightarrow_y^{G_w} C_w$, and C_w has maximum size among all such reachable sets. Define $C_u = C_v \cup C_w$. We will now show that C_u is reachable from both $A \cap V_u$ and $B \cap V_u$, which proves Claim A for node u.

Lemma 13 shows that there exists an independent set J of G_u with $(A \cap V_u) \leftrightarrow_\ell^{G_u} J$ and $|J \cap V_v| = x$. Using this, we argue that there exists an independent set J_1 of G_u with $(A \cap V_u) \leftrightarrow_\ell^{G_u} J_1$ and $J_1 \cap V_v = C_v$. If $A \cap V_u = \emptyset$, then this claim is trivial. Otherwise, we can apply (module) Lemma 7 to draw this conclusion (using V_v, G_u, J and C_v in the roles of the module M, entire graph G, and independent sets B and C, respectively). Analogously, we may conclude that there exists an independent set J_2 of G_u with $(A \cap V_u) \leftrightarrow_\ell^{G_u} J_2$ and $J_2 \cap V_w = C_w$. Since $C_u = C_v \cup C_w$, we can now apply (module) Lemma 8 (with G_u in the role of the entire graph, V_v and V_w in the roles of disjoint modules M_1 and M_2, and J_1 and J_2 in the roles of B_1 and B_2), to conclude that $A \leftrightarrow_\ell^{G_u} C_u$. For this, we require the fact that C_v has maximum size among all independent sets of G_v that are reachable from $A \cap V_v$.

The argument from the previous paragraph also holds when replacing A by B, since C_v and C_w are also maximum reachable independent sets from $B \cap V_v$ and $B \cap V_w$. Thus we may also conclude that $B \leftrightarrow_\ell^{G_u} C_u$. Using the fact that $\leftrightarrow_\ell^{G_u}$ is an equivalence relation, we conclude that $A \leftrightarrow_\ell^{G_u} B$, which proves the desired claim for u.

This concludes the induction proof of Claim A. Applying Claim A to the root node r of T shows that $A \leftrightarrow_k^G B$, since $\lambda_k(r) = k$ (Proposition 25), and $G = G_r$, and therefore concludes the proof of the theorem. □

Next, we prove our main algorithmic result. In the next section, we give examples of graph classes for which this theorem yields efficient algorithms.

Theorem 5 *Let T be a generalized cotree for a graph G on n vertices, let $k \in \mathbb{N}$ and let A and B be independent sets of G. If for every nontrivial leaf $v \in V(T)$ and relevant integer ℓ,*

- *the values $\mu_\ell^A(v)$ and $\mu_\ell^B(v)$ are known, and*
- *it is known whether $(A \cap V_v) \leftrightarrow_\ell^{G_v} (B \cap V_v)$,*

then in time $O(n^2)$ it can be decided whether $A \leftrightarrow_k^G B$.

Proof: We may assume that $|A| \geq k$ and $|B| \geq k$, otherwise we can immediately answer NO. First we use a *bottom up* dynamic programming algorithm, to compute the values $\mu_\ell^A(u)$ and $\mu_\ell^B(u)$ for every node u and relevant integer ℓ. Theorem 24 shows that this can be done in time $O(n^2)$, and that at the same time the maximum ℓ-stable tuples can be computed for A, B, all union nodes u and relevant integers ℓ. (Recall that this uses the dynamic programming rules for trivial leaves, join nodes and union nodes given in Proposition 10, Proposition 11 and Lemma 19, respectively, and the fast computation of maximum ℓ-stable tuples given in Section 5.3.)

Next, we start the *top down* phase of the dynamic programming algorithm, where we compute the values $\lambda_k^A(u)$ and $\lambda_k^B(u)$ for every node u. For the root node r of T, we can initialize these values to k (Proposition 25). Next, for every node u for which these two values are known, we can compute these two values for the two children v and w, by applying Proposition 26 for join nodes and Lemma 27 for union nodes. Note that applying Lemma 27 to a union node u requires the previously computed maximum ℓ-stable tuple (x, y) for $I = A, B$, with $\ell = \lambda_k^I(u)$. This is why the bottom up phase is required.

Finally, we return YES if

- for all nodes $v \in V(T)$, $\lambda_k^A(v) = \lambda_k^B(v)$, and
- for all leaves $v \in V(T)$, $(A \cap V_v) \leftrightarrow_\ell^{G_v} (B \cap V_v)$, where $\ell = \lambda_k^A(v)$.

This is correct by Theorem 3. (Note that for trivial leaves $v \in V(T)$, $(A \cap V_v) \leftrightarrow_\ell^{G_v} (B \cap V_v)$ always holds, and for nontrivial leaves, we assume that this information is given.) Considering the dynamic programming rules, every value that is assigned in the top down phase can be computed in constant time per value. Hence the top down phase takes time $O(|V(T)|) = O(n)$, and the total complexity of the algorithm becomes $O(n^2)$ (which is dominated by the bottom up phase). □

7 Graph Classes

In this section, we discuss graph classes to which Theorem 5 applies. Firstly, Theorem 5 easily implies that the TAR-Reachability problem can be decided efficiently on cographs.

Theorem 28 *Let G be a cograph on n vertices, let $k \in \mathbb{N}$ and let A and B be independent sets of G. In time $O(n^2)$ it can be decided whether $A \leftrightarrow_k^G B$.*

Proof: A cotree T for G can be constructed in linear time [9]. We can easily guarantee that this is a binary tree. Since a cotree only has trivial leaves, Theorem 5 can now be applied. Indeed, for a trivial leaf u: Proposition 10 shows that $\mu_\ell^I(u) = 1$ holds for all relevant $\ell \in \{0, 1\}$. Secondly, it can be seen that $(A \cap V_v) \leftrightarrow_\ell^{G_v} (B \cap V_v)$ always holds for all relevant $\ell \in \{0, 1\}$. So the conditions of Theorem 5 are satisfied. □

Combining this theorem with Lemma 1 shows that we can efficiently decide whether $A \leftrightarrow_{\mathrm{TJ}}^G B$ in the case that G is a cograph, which answers an open question from [21]:

Corollary 29 *Let G be a cograph on n vertices, and let A and B be independent sets of G. In time $O(n^2)$ it can be decided whether $A \leftrightarrow_{TJ}^G B$.*

Theorem 5 is however much stronger, and implies that TAR-Reachability can be decided efficiently for much richer graph classes. We will now give an example of such a graph class, namely the class of *all graphs that admit a cotree decomposition into chordal graphs.* Along the way, we will introduce some tools that allow proving the same for other graph classes that can be obtained by taking unions and joins of graphs from a graph class $\mathcal{G}$, for which the values/properties from Theorem 5 can efficiently be computed/decided.

A graph G is *chordal* if it contains no cycles of length four or more as induced subgraphs (in other words: if every cycle of length at least four contains a *chord*). The only two properties of chordal graphs that we will use are the following. Firstly:

Theorem 30 ([14]) *Let G be a chordal graph. Then $\alpha(G)$ can be computed in polynomial time.*

This statement is well-known, and relatively easy to prove using the concept of simplicial vertices. For the more general class of *perfect graphs*, $\alpha(G)$ can in fact also be computed in polynomial time. See [23, Section 66.3] for more background. Secondly, we use the fact that chordal graphs are obviously even-hole-free. A graph G is *even-hole-free* if it contains no even cycles as induced subgraphs. To our knowledge, no polynomial time algorithm for computing $\alpha(G)$ for even-hole-free graphs is known; otherwise, the result from this section could be generalized to even-hole-free graphs. See also [24, 21]. We will also apply the following result, which was proved in [21].

Theorem 31 ([21]) *Let A and B be two independent sets of an even-hole-free graph G with $|A| = |B|$. Then $A \leftrightarrow_{TJ}^G B$.*

Using Lemma 1, this theorem can be applied to the TAR model to conclude:

Lemma 32 *Let A and B be two distinct independent sets of an even-hole-free graph G. Then $A \leftrightarrow_k^G B$ if and only if neither A nor B is a dominating set of size k.*

Proof: If A is a dominating set of size k, then no token can be added to A, and no token can be removed from A. So A has no neighbors in $\mathrm{TAR}_k(G)$. Since A and B are distinct, it follows that $A \not\leftrightarrow_k^G B$. This follows similarly if B is a dominating set of size k.

Now suppose that neither A nor B is a dominating set of size k. Then we show that $A \leftrightarrow_k^G B$. We can easily construct an independent set A' with $A \leftrightarrow_k^G A'$ and $|A'| = k+1$:

- If $|A| = k$ then add an arbitrary vertex v which has no neighbors in A (which exists since A is not dominating).
- If $|A| \geq k+1$ then remove arbitrary vertices from A until an independent set of size $k+1$ is obtained.

Similarly, we can easily construct an independent set B' with $B \leftrightarrow_k^G B'$ and $|B'| = k+1$. By Theorem 31, $A' \leftrightarrow_{\mathrm{TJ}}^G B'$. Next, Lemma 1 shows that $A' \leftrightarrow_k^G B'$. Combining this with $A \leftrightarrow_k^G A'$ and $B \leftrightarrow_k^G B'$ shows that $A \leftrightarrow_k^G B$. □

The above lemma easily yields the following statement.

Corollary 33 *Let I be an independent set of an even-hole-free graph G, and let J be an independent set of G with $I \leftrightarrow_k J$ that maximizes $|J|$ among all such sets. Then*

- *$|J| = |I|$ if I is a dominating set of size k, and*
- *$|J| = \alpha(G)$ otherwise.*

This in turn gives us immediately an easy way to compute the values $\mu_\ell^I(u)$ for the case that G_u is even-hole-free:

Corollary 34 *Let T be a cotree decomposition of a graph G into even-hole-free graphs, and let I be an independent set of G. Then for every leaf $u \in V(T)$, and every relevant value ℓ:*

- *$\mu_\ell^I(u) = |I|$ if $I \cap V_u$ is a dominating set of G_u of size ℓ, and*
- *$\mu_\ell^I(u) = \alpha(G_u)$ otherwise.*

Combining Theorem 30, Lemma 32 and Corollary 34 shows that if we have a cotree decomposition of a graph G into chordal graphs, then the conditions of Theorem 5 are satisfied, so we can compute in polynomial time whether $A \leftrightarrow_k B$. However, it remains to discuss how in general, a cotree decomposition into chordal graphs can be found. Recall that for a graph G, by $\overline{G}$ the *complement* of G is denoted, which is the graph $\overline{G} = (V(G), \{uv \mid uv \notin E(G)\})$.

Definition 35 *A graph H is* indecomposable *if both H and $\overline{H}$ are connected. A* maximal cotree decomposition *of a graph G is a generalized cotree decomposition T such that for every leaf $u \in V(T)$, G_u is indecomposable.*

Proposition 36 *For any graph G, a maximal cotree decomposition of G can be computed in polynomial time.*

Proof: A polynomial time algorithm for testing whether a given graph is indecomposable follows immediately from the definition (quadratic time in fact). Now consider the following algorithm for constructing a maximal cotree decomposition of G: start with a trivial generalized cotree decomposition T, consisting of one (root) node r with $G_r = G$. As long as the current generalized cotree decomposition T contains a leaf $u \in V(T)$ such that G_u is decomposable, partition the vertices of G_u into new sets V_v and V_w such that G_u is the disjoint union or complete join of $G[V_v]$ and $G[V_w]$ (this can be trivially done in the case where G_u is disconnected, respectively in the case where $\overline{G_u}$ is disconnected). Now add corresponding new leaf nodes v and w as children of u, and make u into a union or join node, respectively. This way, a generalized cotree decomposition of G is maintained. The algorithm terminates after at most $|V(G)|$ steps (which all take polynomial time), since in every step, the number of leaves of T increases by one, and a generalized cotree decomposition has at most $|V(G)|$ leaves. When the algorithm terminates, the resulting generalized cotree decomposition is clearly maximal. □

A graph class $\mathcal{G}$ is called *hereditary* if for every $G \in \mathcal{G}$ and every induced subgraph H of G, $H \in \mathcal{G}$ holds. Clearly, chordal graphs are hereditary.

Lemma 37 *Let $\mathcal{G}$ be a hereditary graph class, and let G be a graph that admits a cotree decomposition into $\mathcal{G}$-graphs. Then every maximal cotree decomposition of G is a cotree decomposition into $\mathcal{G}$-graphs.*

Proof: Let T^* be a maximal cotree decomposition of G, and let T^C be a cotree decomposition of G into $\mathcal{G}$-graphs. Denote by G^*_u and G^C_u the subgraphs of G that correspond to nodes $u \in V(T^*)$ and $u \in V(T^C)$, respectively. Similarly, denote their vertex sets by V^*_u and V^C_u.

Consider a leaf $u \in V(T^*)$. We will prove that G^*_u is also part of the graph class $\mathcal{G}$. If there is a leaf node $v \in V(T^C)$ such that $V^*_u \subseteq V^C_v$, then G^*_u is an induced subgraph of G^C_v, so since $\mathcal{G}$ is hereditary, $G^*_u \in \mathcal{G}$ holds.

Now assume that there is no such leaf node in T^C. Then observe that we may consider a node $w \in V(T^C)$ with $V^*_u \subseteq V^C_w$, which has no child nodes that satisfy this property. (In other words: w is a lowest common ancestor of all nodes x with $V^C_x \cap V^*_u \neq \emptyset$.) Let x and y be the two child nodes of w. So by choice of w, V^*_u can be partitioned into two nonempty sets $S = V^*_u \cap V^C_x$ and $T = V^*_u \cap V^C_y$. If w is a join node, then G^*_u can be written as the complete join of $G[S]$ and $G[T]$, so $\overline{G^*}$ is disconnected, contradicting the fact that it is indecomposable. Similarly, if w is a union node, then G^*_u can be written as the disjoin union of $G[S]$ and $G[T]$, so it is disconnected, contradicting the fact that it is indecomposable. This concludes the proof that for every $u \in V(T^*)$, $G_u \in \mathcal{G}$ holds. □

We now summarize how the previous statements yield a polynomial time algorithm for testing $A \leftrightarrow^G_k B$, whenever G is a graph that admits a cotree decomposition into chordal graphs.

Theorem 38 *Let G be a graph that admits a cotree decomposition into chordal graphs, and let A and B be independent sets of G, both of size at least k. Then in polynomial time, we can decide whether $A \leftrightarrow^G_k B$.*

Proof: We first construct a maximal cotree decomposition T of G in polynomial time (Proposition 36). By Lemma 37, T is then a cotree decomposition into chordal graphs (since chordal graphs are hereditary). So by Theorem 30, we can compute $\alpha(G_u)$ for every leaf $u \in V(T)$.

Combining this with Corollary 34, and the fact that chordal graphs are even-hole-free, shows that we can compute the values $\mu_\ell^A(u)$ and $\mu_\ell^B(u)$ for every leaf $u \in V(T)$ and relevant ℓ. Finally, Lemma 32 gives an easy way to decide in polynomial time whether $A \leftrightarrow_\ell^{G_u} B$ for any leaf $u \in V(T)$ and relevant value ℓ. So the conditions of Theorem 5 are satisfied for the generalized cotree decomposition T of G, and thus we can compute in polynomial time whether $A \leftrightarrow_k^G B$. □

8 A Linear Bound on the Diameter of the Solution Graph

Using the previous lemmas, we can efficiently decide whether there exists a k-TAR-sequence from A to B in a cograph G. However, from these lemmas, one cannot easily deduce a polynomial upper bound for the length of such a sequence. This requires studying the aforementioned cascading sequences in more detail, which is what we will do in this section. We will show that if $A \leftrightarrow_k^G B$, then there exists a k-TAR-sequence from A to B of length at most $4n - |A| - |B|$, where $n = |V(G)|$.

The main idea is as follows. Given independent sets A and B of G with $A \leftrightarrow_k^G B$, we choose an appropriate subgraph G' of G such that there exist maximum independent sets A' and B' of G and short k-TAR-sequences from A to A' and from B to B'. These sequences are short in the sense that for every vertex $v \in V(G)$, no token is added on v after the first token is removed from v. So in total, there are at most $2n - |A| - |A'|$ token additions/removals used in the sequence from A to A', and a similar statement holds for B and B'. Finally, we show that k-TAR-sequence from A' to B' exists, of length at most $|A \Delta B|$. (Recall that $A \Delta B = (A \setminus B) \cup (B \setminus A)$ denotes the *symmetric difference* of A and B.) Combining these three k-TAR-sequences yields a k-TAR-sequence from A to B of length at most $4n - |A| - |B|$ in the subgraph G', and therefore also in G.

We now define the type of TAR-sequences that we will consider. For a node $u \in V(T)$ and *every* value of $\ell \in \{0, \ldots, |I \cap V_u|\}$, a *subsequence* of the next sequence outlines an ℓ-TAR-sequence for G_u from $I \cap V_u$ to an independent set J with $|J| = \mu_\ell^I(u)$ (Properties 1 and 4). In addition it is *short* in the sense that every vertex of G_u is used for at most one token addition (Property 3). This motivates the name *short universal sequence*. Examples of these sequences are given in Figure 4 for the graphs G_v and G_w from Figure 3.

Definition 39 *Let T be a cotree of a graph G, and let I be an independent set of G. A* short universal sequence *or* SU-sequence *for a node $u \in V(T)$, based on I, is a sequence $C_0, \ldots, C_p$ of independent sets of G_u that satisfy the following properties:*

1. *$C_0 = I \cap V_u$.*
2. *For all $i \in \{0, \ldots, p-1\}$: $|C_{i+1}| > |C_i|$.*
3. *For all $i, j \in \{0, \ldots, p-1\}$ with $i \leq j$: $(C_{j+1} \setminus C_j) \cap C_i = \emptyset$.*
4. *For all $\ell \in \{0, \ldots, |I \cap V_u|\}$ and $i \in \{0, \ldots, p\}$: if $|C_i| < \mu_\ell^I(u)$ then $i < p$ and there exists an ℓ-TAR-sequence in G_u from C_i to C_{i+1} that only adds tokens on $C_{i+1} \setminus C_i$*

Below we will show by induction over T that for every node $u \in V(T)$, an SU-sequence exists. But first, we will prove two properties that indicate why these sequences are useful for finding short ℓ-TAR-sequences from $I \cap V_u$ to an independent set J with $|J| = \mu_\ell^I(u)$, for any value of ℓ.

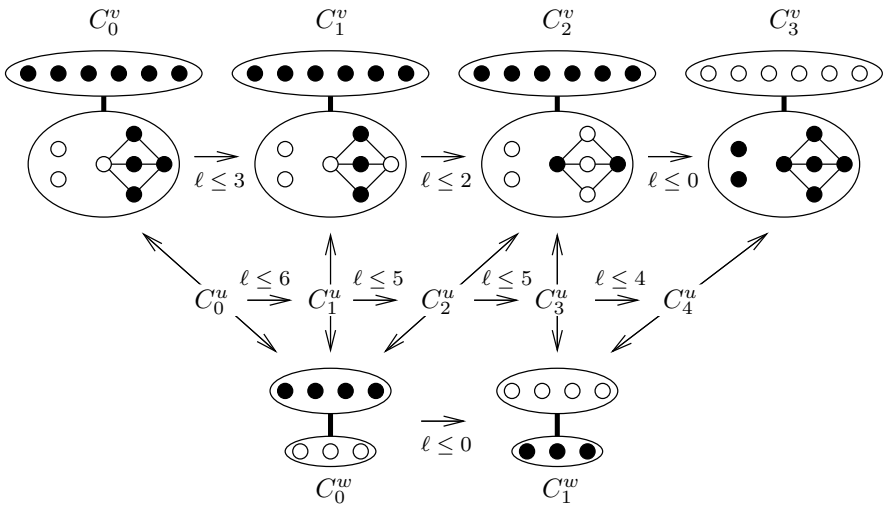

Figure 4: SU-sequences $C_0^v, \ldots, C_3^v$ and C_0^w, C_1^w for the graphs G_v and G_w from Figure 3, and a resulting SU-sequence $C_0^u, \ldots, C_4^u$ for their disjoint union G_u.

Proposition 40 *Let T be a cotree of a graph G, and let I be an independent set of G. Let $C_0, \ldots, C_p$ be an SU-sequence for a node $u \in V(T)$, based on I. Then $|C_p| = \alpha(G_u)$.*

Proof: For all $i \in \{0, \ldots, p\}$, since C_i is an independent set of G_u, $|C_i| \leq \alpha(G_u)$. If the inequality is strict, then $|C_i^u| < \alpha(G_u) = \mu_0^I(u)$, so Property 4 of Definition 39 shows that $i < p$. □

Lemma 41 *Let T be a cotree of a graph G, and let I be an independent set of G. Let $C_0, \ldots, C_p$ be an SU-sequence for a node $u \in V(T)$, based on I. For any $\ell \in \{0, \ldots, |I \cap V_u|\}$ and $i \in \{0, \ldots, p-1\}$ with $|C_i| < \mu_\ell^I(u)$, there exists an ℓ-TAR-sequence from $I \cap V_u$ to C_{i+1} in G_u of length $2(|\bigcup_{j=0}^{i+1} C_j|) - |C_0| - |C_{i+1}|$. Therefore, $\mu_\ell^I(u) \geq |C_{i+1}|$.*

Proof: For any ℓ, we prove the statement by induction over i. We will assume that $|C_i| < \mu_\ell^I(u)$, otherwise there is nothing to prove.

For $i = 0$, Property 4 shows that there is an ℓ-TAR-sequence in G_u from C_0 to C_1 that only adds tokens on $C_1 \backslash C_0$. It follows that this TAR-sequence uses exactly $|C_1 \backslash C_0|$ token additions, and $|C_0 \backslash C_1|$ token removals. We can write

$$|C_1 \backslash C_0| + |C_0 \backslash C_1| = (|C_0 \cup C_1| - |C_0|) + (|C_0 \cup C_1| - |C_1|),$$

which proves the statement.

Now suppose $i \geq 1$. Since $|C_i| < \mu_\ell^I(u)$, Property 2 shows that $|C_{i-1}| < \mu_\ell^I(u)$ holds as well. So by induction, there exists an ℓ-TAR-sequence in G_u from C_0 to C_i of length $2(|\bigcup_{j=0}^{i} C_j|) - |C_0| - |C_i|$. By Property 4, there also exists an ℓ-TAR-sequence in G_u from C_i to C_{i+1}, of length $|C_i \backslash C_{i+1}| + |C_{i+1} \backslash C_i|$ (using an argument similar to above). These can be combined into an ℓ-TAR-sequence in G_u from C_0 to C_{i+1}. The total length of this sequence is therefore:

$$2\left|\bigcup_{j=0}^{i} C_j\right| - |C_0| - |C_i| + |C_i \backslash C_{i+1}| + |C_{i+1} \backslash C_i| =$$

$$2\Big|\bigcup_{j=0}^{i} C_j\Big| - |C_0| - |C_i| + 2|C_i \cup C_{i+1}| - |C_i| - |C_{i+1}| =$$

$$2\Big|\bigcup_{j=0}^{i} C_j\Big| + 2|C_{i+1}\backslash C_i| - |C_0| - |C_{i+1}| \overset{\text{(Property 3)}}{=}$$

$$2\Big|\bigcup_{j=0}^{i} C_j\Big| + 2\Big|C_{i+1}\backslash\Big(\bigcup_{j=0}^{i} C_j\Big)\Big| - |C_0| - |C_{i+1}| =$$

$$2\Big|\bigcup_{j=0}^{i+1} C_j\Big| - |C_0| - |C_{i+1}|.$$

This concludes the induction proof, so we conclude that for any i with $|C_i| < \mu_\ell^I(u)$, there exists an ℓ-TAR-sequence from $I \cap V_u$ to C_{i+1} in G_u of length $2(|\bigcup_{j=0}^{i+1} C_j|) - |C_0| - |C_{i+1}|$. From this, it follows immediately that $\mu_\ell^I(u) \geq |C_{i+1}|$. □

We will now prove that SU-sequences exist for every $u \in V(T)$.

Proposition 42 *Let T be a cotree of a graph G, and let I be an independent set of G. For every leaf node $u \in V(T)$, there exists an SU-sequence based on I.*

Proof: We define $C_0 = I \cap V_u$ and $C_p = \{u\}$. Choose $p = 0$ if these two sets are the same, and $p = 1$ otherwise. One can easily verify that the four properties from Definition 39 hold for this sequence. (Recall that by Proposition 10, $\mu_\ell^I(u) = 1$ for all ℓ.) □

Lemma 43 *Let T be a cotree of a graph G, and let I be an independent set of G. Let $u \in V(T)$ be a join node with children v and w. If there exist SU-sequences for v and w, then there exists an SU-sequence for u (all based on I).*

Proof: Suppose u is a join node, with children v and w. We will construct an SU-sequence $C_0^u, \ldots, C_{p^u}^u$ for u.

If $I \cap V_u = \emptyset$, then we choose $p^u = 1$, $C_0^u = \emptyset$ and C_1^u to be a maximum independent set of G_u. This choice satisfies the four properties of Definition 39. So now we may assume w.l.o.g. that $I \cap V_v \neq \emptyset$ and $I \cap V_w = \emptyset$. By induction, for v there exists an SU-sequence $C_0^v, \ldots, C_{p^v}^v$ based on I. For all $i \in \{0, \ldots, p^v\}$, we choose $C_i^u = C_i^v$. If C_i^v is also a maximum independent set of G_u then this is the entire sequence for u (so we choose $p^u = p^v$), and it satisfies the four properties again (recall that in this case, $\mu_\ell^I(u) = \mu_\ell^I(v)$ for all ℓ, by Proposition 11). Otherwise, since u is a join node, any set $J \subseteq V_u$ is a maximum independent set for G_u if and only if it is a maximum independent set for G_v. Therefore we can choose $p^u = p^v + 1$, and choose $C_{p^u}^u$ to be any maximum independent set of G_v. Clearly, there exists a 0-TAR-sequence from $C_{p^u-1}^u = C_{p^v}^v$ to $C_{p^u}^u$ that only adds tokens on $C_{p^u}^u \backslash C_{p^u-1}^u$. Since $\mu_0^I(u) = |C_{p^u}^u|$ and $\mu_\ell^I(u) = \mu_\ell^I(v)$ for all $\ell \geq 1$ (Proposition 11), this shows that Property 4 again holds for the new sequence. For all $i < p^u$, $C_i^u \subseteq V_v$, so $C_{p^u}^u \cap C_i^u = \emptyset$, and thus Property 3 is again satisfied for this new sequence. Property 1 holds since $C_0^u = |I \cap V_v| = |I \cap V_u|$ (using induction, and that u is a join node with $I \cap V_v \neq \emptyset$, respectively). □

The proof of the following lemma is also illustrated in Figure 4. Given SU-sequences for children v and w of a union node u, we obtain an SU-sequence for u by letting every set in

the new sequence be the union of one set from the SU-sequence for v and one set from the SU-sequence for w.

Lemma 44 *Let T be a cotree of a graph G, and let I be an independent set of G. Let $u \in V(T)$ be a union node with children v and w. If there exist SU-sequences for v and w, then there exists an SU-sequence for u (all based on I).*

Proof: Let $C_0^v, \ldots, C_{p^v}^v$ and $C_0^w, \ldots, C_{p^w}^w$ be SU-sequences based on I for v and w, respectively. We will construct a SU-sequence $C_0^u, \ldots, C_{p^u}^u$ for u from these. These sets will be constructed such that for every index a, there exist indices b and c with

$$C_a^u = C_b^v \cup C_c^w.$$

First we choose $C_0^u = C_0^v \cup C_0^w$, which guarantees that $C_0^u = I \cap V_u$, so Property 1 is satisfied. Next, for every choice of indices a, b, c such that we assigned $C_a^u = C_b^v \cup C_c^w$, continue the construction of the sequence according to the following method. For notational convenience, we define $\mu_i^I(v) = \mu_0^I(v)$ and $\mu_i^I(w) = \mu_0^I(w)$ for all $i < 0$.

(a) Denote $q = |C_b^v|$ and $r = |C_c^w|$.

(b) If $q = \alpha(G_v)$ and $r = \alpha(G_w)$ then assign $p^u := a$ (so the SU-sequence for u ends here). Otherwise, choose C_{a+1}^u as follows:

(c) Choose ℓ to be the maximum value in $\{0, \ldots, |I \cap V_u|\}$ such that $\mu_{\ell-r}^I(v) > q$ or $\mu_{\ell-q}^I(w) > r$.

(d) If $\mu_{\ell-r}^I(v) > q$ then choose $C_{a+1}^u = C_{b+1}^v \cup C_c^w$, and otherwise (when $\mu_{\ell-q}^I(w) > r$) choose $C_{a+1}^u = C_b^v \cup C_{c+1}^w$.

We first argue that a value ℓ can always be chosen as in (c): If $q = |C_b^v| < \alpha(G_v)$, then by Property 4, $\mu_0^I(v) > q$. So choosing any $\ell \in \{0, \ldots, |I \cap V_u|\}$ with $\ell \leq r$ suffices. Otherwise, by (b), $r < \alpha(G_w)$, and any ℓ with $\ell \leq q$ suffices by an analog argument. The above construction defines the sequence $C_0^u, \ldots, C_{p^u}^u$. We will now prove that it is an SU-sequence.

As observed above, $C_0^u = C_0^v \cup C_0^w = (I \cap V_v) \cup (I \cap V_w) = I \cap V_u$, so Property 1 is satisfied. Since $|C_{b+1}| > |C_b|$ and $|C_{c+1}| > |C_c|$ holds for any b and c (Property 2), it follows that $|C_{a+1}| > |C_a|$ holds for any $a < p^u$, which proves Property 2 for the new sequence.

Now we prove Property 3. Consider $a' \leq a$ with $C_a^u = C_b^v \cup C_c^w$ and $C_{a'}^u = C_{b'}^v \cup C_{c'}^w$. So $b' \leq b$ and $c' \leq c$. Assume w.l.o.g. that $C_{a+1}^u = C_{b+1}^v \cup C_c^w$. Then we can write

$$(C_{a+1}^u \backslash C_a^u) \cap C_{a'}^u \subseteq (C_{b+1}^v \backslash C_b^v) \cap (C_{b'}^v \cup C_{c'}^w) = \emptyset.$$

For the last equality, we used

- Property 3 for v to conclude that $(C_{b+1}^v \backslash C_b^v) \cap C_{b'}^v = \emptyset$, and
- the observations that $C_{b+1}^v \subseteq V_v$, $C_{c'}^w \subseteq V_w$, and $V_v \cap V_w = \emptyset$ to conclude that $(C_{b+1}^v \backslash C_b^v) \cap C_{c'}^w = \emptyset$.

It remains to prove Property 4 for the new sequence. First note that $\mu_0^I(u) = \alpha(G_u) = \alpha(G_v) + \alpha(G_w)$, so we may end the sequence when $q = \alpha(G_v)$ and $r = \alpha(G_w)$. Now consider any index $a < p^u$, such that we constructed C_{a+1}^u from $C_a^u = C_b^v \cup C_c^w$ using the above

method. We will prove for all $\ell' \in \{0, \ldots, |I \cap V_u|\}$ with $|C_a| < \mu^I_{\ell'}(u)$ that there exists an ℓ'-TAR-sequence in G_u from C_a to C_{a+1} that only adds tokens on $C_{a+1} \setminus C_a$.

First, we show that $|C^u_a| < \mu^I_{\ell'}(u)$ implies that $\mu^I_{\ell'-r}(v) > q$ or $\mu^I_{\ell'-q}(w) > r$. Consider an ℓ'-TAR-sequence $I_0, \ldots, I_q$ from $I \cap V_u$ to an independent set J of G_u with $|J| = \mu^I_{\ell'}(u) > |C^u_a| = q + r$. Let j be the first index such that $|I_j \cap V_v| \geq q + 1$ or $|I_j \cap V_w| \geq r + 1$. Clearly, such an index j exists, and $j \geq 1$ holds since $|C^u_a| \geq |C^u_0|$ by Property 2. W.l.o.g. assume that $|I_j \cap V_v| \geq q + 1$. Then define $I'_i = I_i \cap V_v$ for all $i \in \{0, \ldots, j\}$. By choice of i, for all $i \in \{0, \ldots, j\}$ it holds that $|I_i \cap V_w| \leq r$, and therefore $|I'_i| \geq \ell' - r$. So the sequence $I'_0, \ldots, I'_j$ is an $(\ell' - r)$-TAR-sequence for G_v from $I \cap V_v$ to an independent set I'_j with $|I'_j| \geq q + 1$, and thus $\mu^I_{\ell'-r}(v) \geq q + 1$.

From this fact we conclude that for any $\ell' \in \{0, \ldots, |I \cap V_u|\}$ with $|C_a| < \mu^I_{\ell'}(u)$, it holds that $\ell' \leq \ell$, where ℓ is the value chosen in (c). We conclude the proof of Property 4 by showing that there exists an ℓ-TAR-sequence in G_u from C^u_a to C^u_{a+1} that only adds tokens on $C^u_{a+1} \setminus C^u_a$ (which is then obviously also an ℓ'-TAR-sequence).

Consider the case that we have chosen $C^u_{a+1} = C^v_{b+1} \cup C^w_c$. Then $\mu^I_{\ell-r}(v) > |C^v_b|$, so by using Property 4 for the SU-sequence for v, there exists an $(\ell - r)$-TAR-sequence in G_v from C^v_b to C^v_{b+1} that only adds tokens on $C^v_{b+1} \setminus C^v_b$. If we apply the same token additions to $C^u_a = C^v_b \cup C^w_c$, then this yields the desired ℓ-TAR-sequence from C^u_a to C^u_{a+1}, since any independent set in this sequence contains r vertices of V_w. If $C^u_{a+1} = C^v_b \cup C^w_{c+1}$, then $\mu^I_{\ell-q}(w) > |C^w_c|$, and the proof is analog.

Summarizing, we have now shown that for the constructed sequence $C^u_0, \ldots, C^u_{p^u}$, all properties from Definition 39 hold, and therefore it is an SU-sequence for u, based on I, which concludes the proof of the lemma. □

A straightforward induction proof based on Proposition 42, Lemma 43 and Lemma 44 now yields the following statement.

Theorem 45 *Let T be a cotree of a graph G, and let I be an independent set of G. For every node $u \in V(T)$, there exists an SU-sequence based on I.*

Combined with Proposition 40 and Lemma 41, this shows that for any value of k such that there exists a k-TAR-sequence in G from I to some maximum independent set of G, then there exists a short k-TAR-sequence of this type.

Theorem 46 *Let G be a graph on n vertices, let T be a cotree of G with root r, let I be an independent set of G, and let k be an integer such that $\mu^I_k(r) = \alpha(G)$. Then there exists a k-TAR-sequence from I to some maximum independent set J of G with length at most $2n - |I| - \alpha(G)$.*

Proof: By Theorem 45, there exists an SU-sequence $C_0, \ldots, C_p$ for the root node r. Since $\mu^I_k(r) = \alpha(G) = |C_p|$ (Proposition 40), Lemma 41 shows that there exists a k-TAR-sequence from I to C_p of length $2(|\bigcup_{j=0}^{p} C_j|) - |C_0| - |C_p| \leq 2n - |I| - \alpha(G)$. □

Theorem 46 can be used to prove the existence of a linear length TAR-sequence between any two independent sets A and B with $A \leftrightarrow^G_k B$, by reconfiguring both to a common reachable maximum independent set. There are however two problems with this approach: first, even though $A \leftrightarrow^G_k B$ holds, it may be that A and B cannot reach any maximum independent set of G. This is remedied by considering an appropriate subgraph G' of G, such that $A \leftrightarrow^{G'}_k B$ holds, and both A and B can reach a maximum independent set of

G'. Lemma 48 below indicates how this graph G' can be chosen – it suffices to simply omit all vertices that are not in any independent set that can be reached from A or B. Secondly, even if both A and B can both reach a maximum independent set of a graph G (i.e. $\mu_k^A(r) = \alpha(G) = \mu_k^B(r)$), it may be that G has multiple maximum independent sets, and Theorem 46 does not specify which one is reachable. In fact, from the construction of the SU-sequences it can be seen that different choices of A and B may lead to different maximum independent sets. Therefore, to conclude the proof, we also need to demonstrate that short TAR-sequences exist between any pair of maximum independent sets that can reach each other. This is done in the next lemma.

Lemma 47 *Let A and B be two maximum independent sets of a cograph G. If $A \leftrightarrow_k^G B$, then there exists a k-TAR-sequence from A to B of length $|A \Delta B|$.*

Proof: We prove the statement by induction over $|A \Delta B|$. Let T be a cotree of G. If $A = B$ then there is nothing to prove, so assume now that $|A \Delta B| \geq 1$. Define a *difference node* to be a node $u \in V(T)$ with $A \cap V_u = \emptyset$ and $B \cap V_u \neq \emptyset$ or with $A \cap V_u \neq \emptyset$ and $B \cap V_u = \emptyset$.

Consider a *join node u with children v and w such that u is not a difference node, but v and w are*. We first argue that such a node exists. Since $A \backslash B \neq \emptyset$, there exists at least one difference node (a leaf of T). Considering the root r, there exists also at least one node that is not a difference node. So we may consider a difference node v for which the parent u is not a difference node. W.l.o.g. assume that $A \cap V_v \neq \emptyset$ and $B \cap V_v = \emptyset$. Since u is not a difference node, $B \cap V_w \neq \emptyset$, where w is the other child of u. If u is a union node, then we can add $A \cap V_v$ to B, such that the result is a larger independent set (since V_u is a module with $V_u \cap B \neq \emptyset$, and vertices in $A \cap V_v$ are not adjacent to vertices in $B \cap V_w$), a contradiction with the maximality of B. So v is a join node. Since A is an independent set, it follows that $A \cap V_w = \emptyset$, and therefore w is also a difference node, which proves that a node u with the stated properties exists.

Next, we prove that $|A \backslash V_u| \geq k$ and $|B \backslash V_u| \geq k$. Consider a k-TAR-sequence $I_0, \ldots, I_p$ from A to B. By choice of u, this sequence contains an independent set that contains no vertices of V_u. Let I_i be the first such independent set in the sequence. So $i \geq 1$ and $I_{i-1} = I_i \cup \{x\}$ for some $x \in V_u$. Because I_{i-1} is an independent set and V_u is a module, I_i contains no vertices that are adjacent to any vertex in V_u. So $(A \cap V_u) \cup I_i$ is an independent set, which implies that $|A| = \alpha(G) \geq |A \cap V_u| + |I_i| \geq |A \cap V_u| + k$, and thus $|A \backslash V_u| \geq k$. Analogously, $|B \backslash V_u| \geq k$ follows.

Since V_u is a module of G, it follows that $(A \backslash V_u) \cup (B \cap V_u)$ and $(B \backslash V_u) \cup (A \cap V_u)$ are also independent sets. In fact, since their cardinalities sum to $2\alpha(G)$, and neither set can be larger than $\alpha(G)$, it follows that both are maximum independent sets of G. Denote $A' = (A \backslash V_u) \cup (B \cap V_u)$.

Since $|A \backslash V_u| \geq k$, a k-TAR-sequence from A to A' can be obtained by first removing all tokens from $A \cap V_u$, and next adding tokens on all of $B \cap V_u$. This sequence has length $|A \Delta A'|$. By induction, there is a k-TAR-sequence from A' to B of length $|A' \Delta B|$. Because $|A \Delta A'| + |A' \Delta B| = |A \Delta B|$, this proves the statement. □

Lemma 48 *Let T be a cotree for G, and let I be an independent set of G such that for all $v \in V(G)$, there exists an independent set J with $I \leftrightarrow_k^G J$ and $v \in J$. Then $\mu_k^I = \alpha(G)$.*

Proof: Let I^* be a maximum independent set of G. By induction over the cotree T, we will prove that Claim A below holds for every node $u \in V(T)$. Applying Claim A for to the root node of T proves the lemma statement.

Claim A: There exists an independent set J of G with $I^* \cap V_u \subseteq J \cap V_u$ and $I \leftrightarrow_k^G J$.

Suppose $u \in V(T)$ is a (trivial) leaf node. Then Claim A follows immediately from the assumption.

Suppose u is a join node, with children v and w. We may assume w.l.o.g. that $I^* \cap V_w = \emptyset$. By induction, there exists an independent set J with $I \leftrightarrow_k^G J$ and $I^* \cap V_v \subseteq J \cap V_v$. Therefore, $I^* \cap V_u = I^* \cap V_v \subseteq J \cap V_v \subseteq J \cap V_u$, which proves Claim A for u.

Finally, suppose u is a union node, with children v and w. If $I^* \cap V_u = \emptyset$ then Claim A follows trivially for u, so assume this is not the case. Since I^* is now a maximum independent set of G with $I^* \cap V_u \neq \emptyset$, and V_u is a module of G that is the disjoint union of V_v and V_w, it follows that $I^* \cap V_v$ and $I^* \cap V_w$ are maximum independent sets for G_v and G_w, respectively. Indeed, if this would not be the case, then the size of I^* can be increased by replacing $I^* \cap V_v$ or $I^* \cap V_w$ by arbitrary maximum independent sets of G_v and G_w respectively, while maintaining an independent set, a contradiction.

By induction, there exists an independent set J_v with $I \leftrightarrow_k^G J_v$ and $I^* \cap V_v \subseteq J_v \cap V_v$, and there exists an independent set J_w with $I \leftrightarrow_k^G J_w$ and $I^* \cap V_w \subseteq J_w \cap V_w$. It follows that $J_v \cap V_v$ and $J_w \cap V_w$ are maximum independent sets of G_v and G_w respectively. We may now apply (module) Lemma 8 (with V_u, V_v, V_w, I, J_v and J_w in the roles of M, M_1, M_2, A, B_1, B_2, respectively), to conclude that there exists an independent set J of G with $I \leftrightarrow_k^G J$, $J \cap V_v = J_v \cap V_v$, and $J \cap V_w = J_w \cap V_w$. So $I^* \cap V_u = (I^* \cap V_v) \cup (I^* \cap V_w) \subseteq J \cap V_u$. This proves Claim A for u. □

Now we can prove the main theorem from this section.

Theorem 49 *Let G be a cograph on n vertices, with independent sets A and B such that $A \leftrightarrow_k^G B$. Then there exists a k-TAR-sequence from A to B of length at most $4n - |A| - |B|$.*

Proof: For an independent set I of G, call a vertex $v \in V(G)$ *k-accessible from I* if there exists an independent set J with $I \leftrightarrow_k^G J$ and $v \in J$. Since $A \leftrightarrow_k^G B$, and $\leftrightarrow_k^G$ is an equivalence relation, it follows that for every vertex $v \in V(G)$, v is k-accessible from A if and only if it is k-accessible from B. So we may consider the subgraph G' induced by all vertices that are k-accessible from A. For any independent set J of G it now holds that $A \leftrightarrow_k^G J$ if and only if $J \subseteq V(G')$ and $A \leftrightarrow_k^{G'} J$, and the same statement holds if we replace A by B.

Since G' is an induced subgraph of G, it is again a cograph, so we may choose T to be a cotree of G', with root r. Denote $n' = |V(G')|$. By definition, G' satisfies the conditions of Lemma 48, for both $I = A$ and $I = B$, so $\mu_k^A(r) = \alpha(G') = \mu_k^B(r)$. Theorem 46 then shows that there exist k-TAR-sequences from A and B to maximum independent sets A' and B' of G' respectively, of length at most $2n' - |A| - |A'|$ and $2n' - |B| - |B'|$. Lemma 47 shows that there exists a k-TAR-sequence from A' to B' of length $|A' \Delta B'|$. Combining these three k-TAR-sequences gives a k-TAR-sequence from A to B in G' of length at most $4n' - |A| - |B| - |A'| - |B'| + |A' \Delta B'| \leq 4n' - |A| - |B| \leq 4n - |A| - |B|$. Since G' is an induced subgraph of G, this is also a k-TAR-sequence for G. □

This immediately yields:

Corollary 50 *For any cograph G on n vertices and integer k, components of $TAR_k(G)$ have diameter at most $4n - 2k$.*

Combining the previous corollary with Lemma 1 yields:

Corollary 51 *For any cograph G on n vertices and integer k, components of $TJ_k(G)$ have diameter at most $2n - k$.*

9 Discussion

In this paper, we showed that the TAR-Reachability problem (and thus the TJ-Reachability problem) can be solved efficiently for any graph that admits a cograph decomposition into graphs that satisfy certain properties (Theorem 5) – call this a *good graph class*. Chordal graphs are given as an example of a good graph class. In fact, this might be generalized to even-hole-free graphs, provided that the following question can be answered affirmatively: can $\alpha(G)$ be computed in polynomial time if G is an even-hole-free graph? This is a well-known open question [24], and also a negative answer (i.e. NP-hardness proof) would be interesting (see [21]).

Another good graph class is the class of claw-free graphs, which will be shown in another paper [5]. Finally, Theorem 5 easily applies to any graph class such that graphs on n vertices admit a cograph decomposition into $O(\log n)$ sized graphs: in this case, a trivial (exponential time) exhaustive search procedure can be applied to the base graphs, such that the total complexity is still polynomial in n.

Together, this shows that the TAR-Reachability problem can be solved efficiently for quite a rich graph class. Considering the fact that TAR-Reachability is PSPACE-hard for perfect graphs [21], the boundary between hard and easy graph classes for this problem starts to become clear.

Recall that cographs are exactly the graphs of cliquewidth two, and of modular width two [11]. Generalizing our result to an efficient algorithm for graphs of bounded cliquewidth may be too challenging; a more reasonable goal is to first consider graphs of bounded modular width. The *modular width* of a graph is the largest number of vertices of a prime graph appearing at some node of its unique modular decomposition tree [11, 13]. Is there a polynomial time algorithm for TAR-Reachability for all graphs of modular width at most k, for every constant k?

The following two questions related to independent set reconfiguration in cographs are still open: first, what is the complexity of deciding whether there exists a k-TAR-sequence of length at most ℓ between two independent sets of a cograph? (Recall that for general graphs, this is strongly NP-hard [21].) Secondly, what is the complexity of deciding whether $\mathrm{TAR}_k(G)$ is connected, if G is a cograph? We expect that a variant of our DP algorithm can be used to show that this problem can be decided in polynomial time.

References

[1] H. L. Bodlaender. A tourist guide through treewidth. *Acta Cybernetica*, 11:1–23, 1993.

[2] P. Bonsma. Rerouting shortest paths in planar graphs. In *FSTTCS 2012*, volume 18 of *LIPIcs*, pages 337–349. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.

[3] P. Bonsma. The complexity of rerouting shortest paths. *Theoretical Computer Science*, 510:1 – 12, 2013.

[4] P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, 2009.

[5] P. Bonsma, M. Kamiński, and M. Wrochna. Independent set reconfiguration in claw-free graphs. working paper, 2014.

[6] L. Cereceda, J. van den Heuvel, and M. Johnson. Connectedness of the graph of vertex-colourings. *Discrete Applied Mathematics*, 308(5–6):913–919, 2008.

[7] L. Cereceda, J. van den Heuvel, and M. Johnson. Mixing 3-colourings in bipartite graphs. *European Journal of Combinatorics*, 30(7):1593–1606, 2009.

[8] L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *Journal of Graph Theory*, 67(1):69–82, 2011.

[9] D. Corneil, Y. Perl, and L. Stewart. A linear recognition algorithm for cographs. *SIAM Journal on Computing*, 14(4):926–934, 1985.

[10] B. Courcelle, J.A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems*, 33(2):125–150, 2000.

[11] B. Courcelle and S. Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(13):77 – 114, 2000.

[12] C.E.J. Eggermont and G.J. Woeginger. Motion planning with pulley, rope, and baskets. In *STACS 2012*, volume 14 of *LIPIcs*, pages 374–383. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2012.

[13] J. Gajarský, M. Lampis, and S. Ordyniak. Parameterized algorithms for modular-width. In *IPEC 2013*, volume 8246 of *LNCS*, pages 163–176. Springer, 2013.

[14] F. Gavril. Algorithms for minimum coloring, maximum clique, minimum covering by cliques, and maximum independent set of a chordal graph. *SIAM Journal on Computing*, 1(2):180–187, 1972.

[15] P. Gopalan, P.G. Kolaitis, E. Maneva, and C.H. Papadimitriou. The connectivity of boolean satisfiability: Computational and structural dichotomies. *SIAM Journal on Computing*, 38(6), 2009.

[16] R.A. Hearn and E.D. Demaine. PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. *Theoretical Computer Science*, 343(1–2):72–96, 2005.

[17] J. van den Heuvel. The complexity of change. *Surveys in Combinatorics 2013*, pages 127–160, 2013.

[18] T. Ito and E.D. Demaine. Approximability of the subset sum reconfiguration problem. In *TAMC 2011*, volume 6648 of *LNCS*, pages 58–69. Springer, 2011.

[19] T. Ito, E.D. Demaine, N.J.A. Harvey, C.H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12–14):1054–1065, 2011.

[20] M. Kamiński, P. Medvedev, and M. Milanič. Shortest paths between shortest paths. *Theoretical Computer Science*, 412(39):5205–5210, 2011.

[21] M. Kamiński, P. Medvedev, and M. Milanič. Complexity of independent set reconfigurability problems. *Theoretical Computer Science*, 439:9–15, 2012.

[22] A.E. Mouawad, N. Nishimura, V. Raman, N. Simjour, and A. Suzuki. On the parameterized complexity of reconfiguration problems. In *IPEC 2013*, volume 8246 of *LNCS*, pages 281–294. Springer, 2013.

[23] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*, volume 24 of *Algorithms and Combinatorics*. Springer, Berlin, 2003.

[24] K. Vušković. Even-hole-free graphs: a survey. *Applicable Analysis and Discrete Mathematics*, 4(2):219–240, 2010.